



# **ONTAP 9.6 ONTAP REST API reference**

## **ONTAP 9.6 REST API reference**

NetApp  
April 02, 2024

# Table of Contents

- ONTAP 9.6 ONTAP REST API reference . . . . . 1
- REST API reference . . . . . 2
  - Getting started with the ONTAP REST API . . . . . 2
  - Application . . . . . 33
  - Cloud . . . . . 369
  - Cluster . . . . . 406
  - NAS . . . . . 743
  - NVMe . . . . . 1444
  - Name-services . . . . . 1699
  - Networking . . . . . 1827
  - SAN . . . . . 2121
  - SVM . . . . . 2508
  - Security . . . . . 2639
  - SnapMirror . . . . . 2959
  - Storage . . . . . 3073
  - Support . . . . . 3745
- Legal notices . . . . . 3903
  - Copyright . . . . . 3903
  - Trademarks . . . . . 3903
  - Patents . . . . . 3903
  - Privacy policy . . . . . 3903
  - Open source . . . . . 3903

# ONTAP 9.6 ONTAP 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 with ONTAP, as well as details about using each API endpoint. You can learn more about the ONTAP REST API in the ONTAP automation doc site: <https://docs.netapp.com/us-en/ontap-automation/index.html>.

### 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".
- REST API requests will be rejected when the wait queue is full. This is an uncommon occurrence and is designed to prevent a DOS attack.
- 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.
  - Both support concurrent requests. If ONTAP temporarily can't handle additional calls, it will respond with an HTTP error status code of 429 or 503 (depending on the kind of limit hit) and an error message in the body explaining the limit encountered.  
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>[,...]

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

# 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>

```

## Operations on multiple records

### Overview

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

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

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

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

### Rolling back changes on failure

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

### Serial and parallel operations

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

### Asynchronous job and results

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

### Unique identifiers go in the body

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

## Examples

### Creating two volumes

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

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

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

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

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

```

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

```

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

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

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

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

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

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

```

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

```

## Record-based Errors

### ONTAP Error Response Codes

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

### Option 2: Query-based PATCH and DELETE

APIs supporting PATCH or DELETE requests on a resource instance also support 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 request 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
262274	The specified query is either empty or is equivalent to an empty query.	400

Error Code	Description	HTTP Code
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. To use the same list of fields when dealing with multiple versions of ONTAP, specify `ignore_unknown_fields=true`.

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. Braces can be used to specify multiple fields within an object without repeating the object name. Braces can be nested within each other to select individual attributes within an object hierarchy.



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
fields=b.{c,d} // Fetch b.c and b.d

```

## 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 computational 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.
- If the `order_by` value includes the direction (`asc` or `desc`), then you must make sure that the space between the field name and the direction is properly URL encoded when it is sent to the server. You may use either a `%20` or a `+` to encode the space, but if you send a space character, the server will respond with an error (400 Bad Request). Most programming language libraries will automatically do this encoding for you. Check the documentation of your language and/or library.

### 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 error code is an integer returned in a JSON string. The optional `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
7	Resource limit exceeded.	429 or 503
8	Resource in use.	409 or 503
65541	RPC timed out.	500
65552	Generic RPC failure.	500
65562	Internal RPC error	500
262145	Application code returned an unexpected exception.	500
262160	There are too many requests already being processed. Retry after a short delay.	429
262177	Missing value.	400
262179	Unexpected argument. Argument shown in error message body.	400

Error Code	Description	HTTP Code
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
262190	You must provide one or more values to apply your changes.	400
262196	Field cannot be set in this operation.	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
262202	Field is considered secret and should not be provided in the URL.	400
262210	Unable to retrieve all required records within the timeout. This "order_by" query is not supported under the current system load with the current number of records in the collection.	500
262211	POST request on a REST API does not support filtering on an attribute. Attributes must be in the request body.	400
262212	Request is missing required value.	400
262220	Wildcard fields=* is not allowed for CLI-based REST APIs.	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
262249	Field name is not supported for GET requests because it is not a readable attribute.	400
262250	Field name cannot be queried in a GET request because it is not a readable attribute.	400
262254	Invalid JSON input, an array was expected.	400
262255	An array was found in the JSON when it was not expected.	400
262268	The field is not supported as an order_by= field.	400
262277	The query_fields and query parameters may only be specified for GET requests.	400
262282	Property was specified twice.	400
262286	Mismatching braces found in the fields= query.	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.

## Example

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

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

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

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

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

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

Doing a GET on the Location URL provided:

#The call:

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

# The response:

HTTP/1.1 200 OK

Date: Tue, 07 Jun 2022 17:04:50 GMT

Server: libzapid-httpd

X-Content-Type-Options: nosniff

Cache-Control: no-cache,no-store,must-revalidate

Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-ancestors: 'self'

Content-Length: 284

Content-Type: application/hal+json

Vary: Accept-Encoding

```
{
  "owner": {
    "uuid": "001c3044-d15c-11ec-8091-005056bb49e4"
  },
  "name": "test",
  "path": "/api/cluster",
  "access": "readonly",
  "_links": {
    "self": {
      "href": "/api/security/roles/001c3044-d15c-11ec-8091-005056bb49e4/test/privileges/%2Fapi%2Fcluster"
    }
  }
}
```

## 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
- 429 Too Many Requests: Returned when the client has sent more requests than the server can handle. The client should try again later (defined by the Retry-After header)
- 500 Internal Error: Returned for most other internal error codes
- 502 Bad Gateway: Returned if the application is temporarily unreachable. Try again later
- 503 Service Unavailable: Returned if the server is temporarily overloaded. Try again later.

## 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)
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": "voll1",  
      "_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 &rarr; GET /api/private/cli/volume
- volume create &rarr; POST /api/private/cli/volume
- volume modify &rarr; PATCH /api/private/cli/volume
- volume delete &rarr; DELETE /api/private/cli/volume
- volume restrict &rarr; POST /api/private/cli/volume/restrict
- volume show-space &rarr; GET /api/private/cli/volume/space
- volume show-footprint &rarr; GET /api/private/cli/volume/footprint
- cluster add-node &rarr; POST /api/private/cli/cluster/add-node
- cluster add-node-status &rarr; GET /api/private/cli/system/node/add-node-status
- system node coredump show &rarr; GET /api/private/cli/system/node/coredump
- system node coredump delete &rarr; DELETE /api/private/cli/system/node/coredump
- system node coredump delete-all &rarr; 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*, *nfs\_access* 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"
      }
    }
  }
}
```

```

    },
    "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" } }, "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
content-length: 203

```

### 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
```

---

## 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`

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 [DOC 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.local.name	string	query	False	Filter by rpo.components.rpo.local.name
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
statistics.space.used_percent	string	query	False	Filter by statistics.space.used_percent

Name	Type	In	Required	Description
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
statistics.statistics_incomplete	string	query	False	Filter by statistics.statistics_incomplete

Name	Type	In	Required	Description
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
statistics.components.space.reserved_unused	string	query	False	Filter by statistics.components.space.reserved_unused

Name	Type	In	Required	Description
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_incomplete	string	query	False	Filter by statistics.component s.statistics_incomplete
statistics.component s.shared_storage_pool	string	query	False	Filter by statistics.component s.shared_storage_pool
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	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.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">application</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      },
      "snapshots": {
        "href": "/api/resourcelink"
      }
    },
    "creation_timestamp": "string",
    "generation": 0,
    "maxdata_on_san": {
      "app_type": "mongodb",
      "application_components": {
        "file_system": "generic",
        "host_management_url": "string",
        "metadata": {
        },
        "protection_type": {
          "local_rpo": "6_hourly",
          "remote_rpo": "6_hourly"
        },
        "storage_service": {
          "name": "extreme"
        }
      },
      "metadata": {
      },
      "new_igroups": {
        "initiators": {
        },
        "os_type": "aix",
        "protocol": "fc"
      },
      "ocsm_url": "string",
```

```

    "os_type": "aix"
  },
  "mongo_db_on_san": {
    "dataset": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "new_igroups": {
      "initiators": {
      },
      "os_type": "hyper_v",
      "protocol": "fc"
    },
    "os_type": "hyper_v",
    "protection_type": {
      "local_rpo": "hourly",
      "remote_rpo": "none"
    },
    "secondary_igroups": {
    }
  },
  "nas": {
    "application_components": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "cifs_access": {
      "access": "change"
    },
    "nfs_access": {
      "access": "none"
    },
    "protection_type": {
      "local_rpo": "hourly",
      "remote_rpo": "none"
    }
  },
  "oracle_on_nfs": {
    "archive_log": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "db": {

```

```

    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"oracle_on_san": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fc"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "os_type": "aix",
  "protection_type": {

```

```
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"oracle_rac_on_nfs": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "grid_binary": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_crs": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
}
```

```

    }
  },
  "oracle_rac_on_san": {
    "archive_log": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "db": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "db_sids": {
    },
    "grid_binary": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "new_igroups": {
      "initiators": {
      },
      "os_type": "aix",
      "protocol": "fcp"
    },
    "ora_home": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "oracle_crs": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "os_type": "aix",
    "protection_type": {
      "local_rpo": "hourly",
      "remote_rpo": "none"
    },
    "redo_log": {
      "storage_service": {
        "name": "extreme"
      }
    }
  }
}

```

```

},
"protection_granularity": "application",
"rpo": {
  "components": {
    "name": "string",
    "rpo": {
      "local": {
        "description": "string",
        "name": "6_hourly"
      },
      "remote": {
        "description": "string",
        "name": "6_hourly"
      }
    },
    "uuid": "string"
  },
  "local": {
    "description": "string",
    "name": "6_hourly"
  },
  "remote": {
    "description": "string",
    "name": "6_hourly"
  }
},
"san": {
  "application_components": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fc"
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"sql_on_san": {
  "db": {

```

```

    "storage_service": {
      "name": "extreme"
    }
  },
  "log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "hyper_v",
    "protocol": "fc"
  },
  "os_type": "windows",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "temp_db": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"sql_on_smb": {
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "temp_db": {
    "storage_service": {
      "name": "extreme"
    }
  }
}

```



```

},
"state": "creating",
"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
  }
},
"template": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "protocol": "nas",
  "version": 0
},
"uuid": "string",
"vdi_on_nas": {
  "desktops": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"vdi_on_san": {
  "desktops": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "hypervisor": "hyper_v",
  "new_igroups": {
    "initiators": {
    },
    "protocol": "fcp"
  },
},

```

```
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
},
"vsi_on_nas": {
  "datastore": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"vsi_on_san": {
  "datastore": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "hypervisor": "hyper_v",
  "new_igroups": {
    "initiators": {
    },
    "protocol": "fc"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
}
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	
snapshots	<a href="#">href</a>	

metadata

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

protection\_type

Name	Type	Description
local_rpo	string	The local rpo of the application component. Optional in the POST or PATCH body
remote_rpo	string	The remote rpo of the application component. Optional in the POST or PATCH body

storage\_service

Name	Type	Description
name	string	The storage service of the application component. Optional in the POST or PATCH body

maxdata\_on\_san\_application\_components

application-components

Name	Type	Description
file_system	string	Defines the kind of file system that will be installed on this application component. Optional in the POST body
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. Required in the POST body
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. Required in the POST body and optional in the PATCH body
lun_count	integer	The number of LUNs in the application component. Required in the POST body
metadata	array[ <a href="#">metadata</a> ]	
name	string	The name of the application component. Required in the POST body and optional in the PATCH body
protection_type	<a href="#">protection_type</a>	
storage_service	<a href="#">storage_service</a>	

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

#### metadata

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

#### 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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

#### 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. Required in the POST body
application_components	array[ <a href="#">maxdata_on_san_application_components</a> ]	application-components. Optional in the POST or PATCH body
metadata	array[ <a href="#">metadata</a> ]	
new_igroups	array[ <a href="#">maxdata_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
ocsm_url	string	The OnCommand System Manager URL for this application
os_type	string	The name of the host OS running the application. Required in the POST body

#### storage\_service

Name	Type	Description
name	string	The storage service of the database. Optional in the POST or PATCH body

#### 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. Optional in the POST body
replication_factor	integer	The number of data bearing members of the replicaset, including 1 primary and at least 1 secondary. Optional in the POST body



Name	Type	Description
size	integer	The size of the database. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

protection\_type

Name	Type	Description
local_rpo	string	The local rpo of the application. Optional in the POST or PATCH body
remote_rpo	string	The remote rpo of the application. Optional in the POST body

secondary\_igroups

Name	Type	Description
name	string	The name of the initiator group for each secondary. Optional in the POST or PATCH body

## mongo\_db\_on\_san

MongoDB using SAN.

Name	Type	Description
dataset	<a href="#">dataset</a>	
new_igroups	array[ <a href="#">mongo_db_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Optional in the POST body
primary_igroup_name	string	The initiator group for the primary. Required in the POST body and optional in the PATCH body
protection_type	<a href="#">protection_type</a>	
secondary_igroups	array[ <a href="#">secondary_igroups</a> ]	

## application\_components

Name	Type	Description
name	string	The name of the application component. Optional in the POST or PATCH body
share_count	integer	The number of shares in the application component. Optional in the POST body
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member shares. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

## app\_cifs\_access

The list of CIFS access controls.

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

app\_nfs\_access

The list of NFS access controls.

Name	Type	Description
access	string	The NFS access granted. Optional in the POST body
host	string	The name of the NFS entity granted access. Optional in the POST body

nas

A generic NAS application.

Name	Type	Description
application_components	array[ <a href="#">application_components</a> ]	
cifs_access	array[ <a href="#">app_cifs_access</a> ]	The list of CIFS access controls. Optional in the POST body
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the archive log. Optional in the POST or PATCH body

archive\_log

Name	Type	Description
size	integer	The size of the archive log. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

db

Name	Type	Description
size	integer	The size of the database. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the ORACLE_HOME storage volume. Optional in the POST or PATCH body

ora\_home

Name	Type	Description
size	integer	The size of the ORACLE_HOME storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the redo log group. Optional in the POST or PATCH body

redo\_log

Name	Type	Description
mirrored	boolean	Should the redo log group be mirrored? Optional in the POST body
size	integer	The size of the redo log group. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

oracle\_on\_nfs

Oracle using NFS.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
ora_home	<a href="#">ora_home</a>	
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body

Name	Type	Description
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

oracle\_on\_san

Oracle using SAN.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">oracle_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
ora_home	<a href="#">ora_home</a>	
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the Oracle grid binary storage volume. Optional in the POST or PATCH body

grid\_binary

Name	Type	Description
size	integer	The size of the Oracle grid binary storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the Oracle CRS volume. Optional in the POST or PATCH body

oracle\_crs

Name	Type	Description
copies	integer	The number of CRS volumes. Optional in the POST body
size	integer	The size of the Oracle CRS/voting storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST body
storage_service	<a href="#">storage_service</a>	

oracle\_rac\_on\_nfs

Oracle RAC using NFS.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
grid_binary	<a href="#">grid_binary</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
ora_home	<a href="#">ora_home</a>	
oracle_crs	<a href="#">oracle_crs</a>	
protection_type	<a href="#">protection_type</a>	

Name	Type	Description
redo_log	<a href="#">redo_log</a>	

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. Optional in the POST or PATCH body

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

oracle\_rac\_on\_san

Oracle RAC using SAN.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	



Name	Type	Description
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. Optional in the POST or PATCH body
ora_home	ora_home	
oracle_crs	oracle_crs	
os_type	string	The name of the host OS running the application. Required in the POST body
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 schedule
name	string	The local RPO of the component. This indicates how often component snapshots 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	

Name	Type	Description
remote	<a href="#">remote</a>	

#### components

Name	Type	Description
name	string	Component Name
rpo	<a href="#">rpo</a>	
uuid	string	Component UUID

#### local

Name	Type	Description
description	string	A detailed description of the local RPO. This will include details about the snapshot schedule
name	string	The local RPO of the application. This indicates how often application snapshots 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[ <a href="#">components</a> ]	
is_supported	boolean	Is RPO supported for this application? Generation 1 applications did not support snapshots or MetroCluster

Name	Type	Description
local	<a href="#">local</a>	
remote	<a href="#">remote</a>	

#### 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. Optional in the POST or PATCH body
lun_count	integer	The number of LUNs in the application component. Optional in the POST body
name	string	The name of the application component. Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

#### 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. Required in the POST body and optional in the PATCH body

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. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

san

A generic SAN application.

Name	Type	Description
application_components	array[ <a href="#">application_components</a> ]	
new_igroups	array[ <a href="#">san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the db. Optional in the POST or PATCH body

db

Name	Type	Description
size	integer	The size of the db. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the log db. Optional in the POST or PATCH body

log

Name	Type	Description
size	integer	The size of the log db. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

storage\_service

Name	Type	Description
name	string	The storage service of the temp db. Optional in the POST or PATCH body

temp\_db

Name	Type	Description
size	integer	The size of the temp db. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

sql\_on\_san

Microsoft SQL using SAN.

Name	Type	Description
db	<a href="#">db</a>	
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. Required in the POST body and optional in the PATCH body
log	<a href="#">log</a>	
new_igroups	array[ <a href="#">sql_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Optional in the POST body
protection_type	<a href="#">protection_type</a>	
server_cores_count	integer	The number of server cores for the db. Optional in the POST body
temp_db	<a href="#">temp_db</a>	

access

Name	Type	Description
installer	string	SQL installer admin user name. Optional in the POST body

Name	Type	Description
service_account	string	SQL service account user name. Required in the POST body

#### sql\_on\_smb

Microsoft SQL using SMB.

Name	Type	Description
access	<a href="#">access</a>	
db	<a href="#">db</a>	
log	<a href="#">log</a>	
protection_type	<a href="#">protection_type</a>	
server_cores_count	integer	The number of server cores for the db. Optional in the POST body
temp_db	<a href="#">temp_db</a>	

#### 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 snapshots
used	integer	The amount of spacing currently in use by the system to store snapshots

space

Name	Type	Description
available	integer	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 <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>
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 snapshots 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 snapshots
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 snapshots
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	<a href="#">iops</a>	
latency	<a href="#">latency</a>	
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	<a href="#">snapshot</a>	
space	<a href="#">space</a>	

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	<a href="#">storage_service</a>	
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	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 <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

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 snapshots 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 snapshots
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 snapshots
used_percent	integer	The percentage of the originally provisioned space that is currently being used by the application

#### statistics

Name	Type	Description
components	array[ <a href="#">components</a> ]	
iops	<a href="#">iops</a>	
latency	<a href="#">latency</a>	

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	<a href="#">snapshot</a>	
space	<a href="#">space</a>	
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. Optional in the POST body
uuid	string	SVM UUID. Either the SVM name or UUID must be provided to create an application. Optional in the POST body

#### self\_link

Name	Type	Description
self	<a href="#">href</a>	

#### template

Name	Type	Description
_links	<a href="#">self_link</a>	
name	string	The name of the template that was used to provision this application. Optional in the POST body

Name	Type	Description
protocol	string	The protocol access of the template that was used to provision this application
version	integer	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 <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

#### storage\_service

Name	Type	Description
name	string	The storage service of the desktops. Optional in the POST or PATCH body

#### desktops

Name	Type	Description
count	integer	The number of desktops to support. Optional in the POST or PATCH body
size	integer	The size of the desktops. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body
storage_service	<a href="#">storage_service</a>	

#### hyper\_v\_access

Name	Type	Description
service_account	string	Hyper-V service account. Optional in the POST body

## vdi\_on\_nas

A VDI application using NAS.

Name	Type	Description
desktops	<a href="#">desktops</a>	
hyper_v_access	<a href="#">hyper_v_access</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

## 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. Required in the POST body and optional in the PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

## vdi\_on\_san

A VDI application using SAN.

Name	Type	Description
desktops	<a href="#">desktops</a>	
hypervisor	string	The name of the hypervisor hosting the application. Required in the POST body

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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">vdi_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
protection_type	<a href="#">protection_type</a>	

#### storage\_service

Name	Type	Description
name	string	The storage service of the datastore. Optional in the POST or PATCH body

#### datastore

Name	Type	Description
count	integer	The number of datastores to support. Optional in the POST or PATCH body
size	integer	The size of the datastore. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body
storage_service	<a href="#">storage_service</a>	

#### vsi\_on\_nas

A VSI application using NAS.

Name	Type	Description
datastore	<a href="#">datastore</a>	
hyper_v_access	<a href="#">hyper_v_access</a>	

Name	Type	Description
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

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. Required in the POST body and optional in the PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

vsi\_on\_san

A VSI application using SAN.

Name	Type	Description
datastore	<a href="#">datastore</a>	
hypervisor	string	The name of the hypervisor hosting the application. Required in the POST body
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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">vsi_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body



Name	Type	Description
protection_type	<a href="#">protection_type</a>	

application

Applications

Name	Type	Description
_links	<a href="#">_links</a>	
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 snapshots. Support for snapshots was added at generation 2. Any future generation numbers and their feature set will be documented
maxdata_on_san	<a href="#">maxdata_on_san</a>	MAX Data application using SAN.
mongo_db_on_san	<a href="#">mongo_db_on_san</a>	MongoDB using SAN.
name	string	Application Name. This field is user supplied when the application is created. Required in the POST body
nas	<a href="#">nas</a>	A generic NAS application.
oracle_on_nfs	<a href="#">oracle_on_nfs</a>	Oracle using NFS.
oracle_on_san	<a href="#">oracle_on_san</a>	Oracle using SAN.
oracle_rac_on_nfs	<a href="#">oracle_rac_on_nfs</a>	Oracle RAC using NFS.
oracle_rac_on_san	<a href="#">oracle_rac_on_san</a>	Oracle RAC using SAN.

Name	Type	Description
protection_granularity	string	Protection granularity determines the scope of Snapshot operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot operations are performed on the entire application. If the value is "component", Snapshot operations are performed separately on the application components
rpo	<a href="#">rpo</a>	
san	<a href="#">san</a>	A generic SAN application.
sql_on_san	<a href="#">sql_on_san</a>	Microsoft SQL using SAN.
sql_on_smb	<a href="#">sql_on_smb</a>	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	<a href="#">statistics</a>	
svm	<a href="#">svm</a>	
template	<a href="#">template</a>	
uuid	string	Application UUID. This field is generated when the application is created. Required in the URL
vdi_on_nas	<a href="#">vdi_on_nas</a>	A VDI application using NAS.
vdi_on_san	<a href="#">vdi_on_san</a>	A VDI application using SAN.
vsi_on_nas	<a href="#">vsi_on_nas</a>	A VSI application using NAS.
vsi_on_san	<a href="#">vsi_on_san</a>	A VSI application using SAN.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

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...).

- `igroup_name` - All SAN applications require an initiator group to be specified in some way.
- `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, but likely provide more access to the application than is necessary.

- `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*
  - `nfs_access.host` - *0.0.0.0/0*
  - `nfs_access.access` - *rw*
  - `cifs_access.user_or_group` - *everyone*
  - `cifs_access.access` - *full\_access*
  - `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, this is documented in the description of each property. 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](#)
- [DOC Asynchronous operations](#)

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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 snapshots. Support for snapshots was added at generation 2. Any future generation numbers and their feature set will be documented
maxdata_on_san	<a href="#">maxdata_on_san</a>	MAX Data application using SAN.
mongo_db_on_san	<a href="#">mongo_db_on_san</a>	MongoDB using SAN.
name	string	Application Name. This field is user supplied when the application is created. Required in the POST body
nas	<a href="#">nas</a>	A generic NAS application.
oracle_on_nfs	<a href="#">oracle_on_nfs</a>	Oracle using NFS.
oracle_on_san	<a href="#">oracle_on_san</a>	Oracle using SAN.
oracle_rac_on_nfs	<a href="#">oracle_rac_on_nfs</a>	Oracle RAC using NFS.
oracle_rac_on_san	<a href="#">oracle_rac_on_san</a>	Oracle RAC using SAN.
protection_granularity	string	Protection granularity determines the scope of Snapshot operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot operations are performed on the entire application. If the value is "component", Snapshot operations are performed separately on the application components
rpo	<a href="#">rpo</a>	
san	<a href="#">san</a>	A generic SAN application.
sql_on_san	<a href="#">sql_on_san</a>	Microsoft SQL using SAN.

Name	Type	Description
sql_on_smb	<a href="#">sql_on_smb</a>	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	<a href="#">statistics</a>	
svm	<a href="#">svm</a>	
template	<a href="#">template</a>	
uuid	string	Application UUID. This field is generated when the application is created. Required in the URL
vdi_on_nas	<a href="#">vdi_on_nas</a>	A VDI application using NAS.
vdi_on_san	<a href="#">vdi_on_san</a>	A VDI application using SAN.
vsi_on_nas	<a href="#">vsi_on_nas</a>	A VSI application using NAS.
vsi_on_san	<a href="#">vsi_on_san</a>	A VSI application using SAN.



## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    },
    "snapshots": {
      "href": "/api/resourcelink"
    }
  },
  "creation_timestamp": "string",
  "generation": 0,
  "maxdata_on_san": {
    "app_type": "mongodb",
    "application_components": {
      "file_system": "generic",
      "host_management_url": "string",
      "metadata": {
      },
      "protection_type": {
        "local_rpo": "6_hourly",
        "remote_rpo": "6_hourly"
      },
      "storage_service": {
        "name": "extreme"
      }
    },
    "metadata": {
    },
    "new_igroups": {
      "initiators": {
      },
      "os_type": "aix",
      "protocol": "fc"
    },
    "ocsm_url": "string",
    "os_type": "aix"
  },
  "mongo_db_on_san": {
    "dataset": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "new_igroups": {
```

```

    "initiators": {
    },
    "os_type": "hyper_v",
    "protocol": "fc"
  },
  "os_type": "hyper_v",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "secondary_igroups": {
  }
},
"nas": {
  "application_components": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "cifs_access": {
    "access": "change"
  },
  "nfs_access": {
    "access": "none"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"oracle_on_nfs": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "ora_home": {
    "storage_service": {

```

```

    "name": "extreme"
  }
},
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
},
"redo_log": {
  "storage_service": {
    "name": "extreme"
  }
}
},
"oracle_on_san": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fc"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
}
},

```

```

"oracle_rac_on_nfs": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "grid_binary": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_crs": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_rac_on_san": {
    "archive_log": {
      "storage_service": {
        "name": "extreme"
      }
    }
  },
  "db": {

```

```

    "storage_service": {
      "name": "extreme"
    }
  },
  "db_sids": {
  },
  "grid_binary": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fcp"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_crs": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"protection_granularity": "application",
"rpo": {
  "components": {
    "name": "string",
    "rpo": {
      "local": {
        "description": "string",
        "name": "6_hourly"
      }
    }
  }
}

```

```

    },
    "remote": {
        "description": "string",
        "name": "6_hourly"
    }
},
"uuid": "string"
},
"local": {
    "description": "string",
    "name": "6_hourly"
},
"remote": {
    "description": "string",
    "name": "6_hourly"
}
},
"san": {
    "application_components": {
        "storage_service": {
            "name": "extreme"
        }
    },
    "new_igroups": {
        "initiators": {
        },
        "os_type": "aix",
        "protocol": "fcp"
    },
    "os_type": "aix",
    "protection_type": {
        "local_rpo": "hourly",
        "remote_rpo": "none"
    }
},
"sql_on_san": {
    "db": {
        "storage_service": {
            "name": "extreme"
        }
    },
    "log": {
        "storage_service": {
            "name": "extreme"
        }
    }
},

```

```
"new_igroups": {
  "initiators": {
  },
  "os_type": "hyper_v",
  "protocol": "fcp"
},
"os_type": "windows",
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
},
"temp_db": {
  "storage_service": {
    "name": "extreme"
  }
}
},
"sql_on_smb": {
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "temp_db": {
    "storage_service": {
      "name": "extreme"
    }
  }
}
},
"state": "creating",
"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
}

```



```

},
"template": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"protocol": "nas",
"version": 0
},
"uuid": "string",
"vdi_on_nas": {
  "desktops": {
    "storage_service": {
      "name": "extreme"
    }
  },
},
"nfs_access": {
  "access": "none"
},
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
}
},
"vdi_on_san": {
  "desktops": {
    "storage_service": {
      "name": "extreme"
    }
  },
},
"hypervisor": "hyper_v",
"new_igroups": {
  "initiators": {
  },
  "protocol": "fc"
},
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
}
},
"vsi_on_nas": {
  "datastore": {
    "storage_service": {
      "name": "extreme"
    }
  }
}

```

```

    }
  },
  "nfs_access": {
    "access": "none"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"vsi_on_san": {
  "datastore": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "hypervisor": "hyper_v",
  "new_igroups": {
    "initiators": {
    },
    "protocol": "fc"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
}
}
}

```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	
snapshots	<a href="#">href</a>	

metadata

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

protection\_type

Name	Type	Description
local_rpo	string	The local rpo of the application component. Optional in the POST or PATCH body
remote_rpo	string	The remote rpo of the application component. Optional in the POST or PATCH body

storage\_service

Name	Type	Description
name	string	The storage service of the application component. Optional in the POST or PATCH body

maxdata\_on\_san\_application\_components

application-components

Name	Type	Description
file_system	string	Defines the kind of file system that will be installed on this application component. Optional in the POST body
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. Required in the POST body
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. Required in the POST body and optional in the PATCH body
lun_count	integer	The number of LUNs in the application component. Required in the POST body
metadata	array[ <a href="#">metadata</a> ]	
name	string	The name of the application component. Required in the POST body and optional in the PATCH body
protection_type	<a href="#">protection_type</a>	
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body

metadata

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

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

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. Required in the POST body
application_components	array[ <a href="#">maxdata_on_san_application_components</a> ]	application-components. Optional in the POST or PATCH body
metadata	array[ <a href="#">metadata</a> ]	
new_igroups	array[ <a href="#">maxdata_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body

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. Required in the POST body

#### storage\_service

Name	Type	Description
name	string	The storage service of the database. Optional in the POST or PATCH body

#### 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. Optional in the POST body
replication_factor	integer	The number of data bearing members of the replicaset, including 1 primary and at least 1 secondary. Optional in the POST body
size	integer	The size of the database. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

#### mongo\_db\_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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

#### protection\_type

Name	Type	Description
local_rpo	string	The local rpo of the application. Optional in the POST or PATCH body
remote_rpo	string	The remote rpo of the application. Optional in the POST body

#### secondary\_igroups

Name	Type	Description
name	string	The name of the initiator group for each secondary. Optional in the POST or PATCH body

#### mongo\_db\_on\_san

MongoDB using SAN.

Name	Type	Description
dataset	<a href="#">dataset</a>	
new_igroups	array[ <a href="#">mongo_db_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body



Name	Type	Description
os_type	string	The name of the host OS running the application. Optional in the POST body
primary_igroup_name	string	The initiator group for the primary. Required in the POST body and optional in the PATCH body
protection_type	<a href="#">protection_type</a>	
secondary_igroups	array[ <a href="#">secondary_igroups</a> ]	

#### application\_components

Name	Type	Description
name	string	The name of the application component. Optional in the POST or PATCH body
share_count	integer	The number of shares in the application component. Optional in the POST body
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member shares. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

#### app\_cifs\_access

The list of CIFS access controls.

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

#### app\_nfs\_access

The list of NFS access controls.

Name	Type	Description
access	string	The NFS access granted. Optional in the POST body
host	string	The name of the NFS entity granted access. Optional in the POST body

nas

A generic NAS application.

Name	Type	Description
application_components	array[ <a href="#">application_components</a> ]	
cifs_access	array[ <a href="#">app_cifs_access</a> ]	The list of CIFS access controls. Optional in the POST body
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the archive log. Optional in the POST or PATCH body

archive\_log

Name	Type	Description
size	integer	The size of the archive log. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

db

Name	Type	Description
size	integer	The size of the database. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the ORACLE_HOME storage volume. Optional in the POST or PATCH body

ora\_home

Name	Type	Description
size	integer	The size of the ORACLE_HOME storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the redo log group. Optional in the POST or PATCH body

redo\_log

Name	Type	Description
mirrored	boolean	Should the redo log group be mirrored? Optional in the POST body

Name	Type	Description
size	integer	The size of the redo log group. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

oracle\_on\_nfs

Oracle using NFS.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
ora_home	<a href="#">ora_home</a>	
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

oracle\_on\_san

Oracle using SAN.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">oracle_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
ora_home	<a href="#">ora_home</a>	
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the Oracle grid binary storage volume. Optional in the POST or PATCH body

grid\_binary

Name	Type	Description
size	integer	The size of the Oracle grid binary storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

Name	Type	Description
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the Oracle CRS volume. Optional in the POST or PATCH body

oracle\_crs

Name	Type	Description
copies	integer	The number of CRS volumes. Optional in the POST body
size	integer	The size of the Oracle CRS/voting storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST body
storage_service	<a href="#">storage_service</a>	

oracle\_rac\_on\_nfs

Oracle RAC using NFS.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
grid_binary	<a href="#">grid_binary</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
ora_home	<a href="#">ora_home</a>	
oracle_crs	<a href="#">oracle_crs</a>	
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

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. Optional in the POST or PATCH body

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

oracle\_rac\_on\_san

Oracle RAC using SAN.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
db_sids	array[ <a href="#">db_sids</a> ]	
grid_binary	<a href="#">grid_binary</a>	

Name	Type	Description
new_igroups	array[oracle_rac_on_san_new_igroups]	The list of initiator groups to create. Optional in the POST or PATCH body
ora_home	ora_home	
oracle_crs	oracle_crs	
os_type	string	The name of the host OS running the application. Required in the POST body
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 schedule
name	string	The local RPO of the component. This indicates how often component snapshots 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	<a href="#">rpo</a>	
uuid	string	Component UUID

#### local

Name	Type	Description
description	string	A detailed description of the local RPO. This will include details about the snapshot schedule
name	string	The local RPO of the application. This indicates how often application snapshots 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[ <a href="#">components</a> ]	
is_supported	boolean	Is RPO supported for this application? Generation 1 applications did not support snapshots or MetroCluster
local	<a href="#">local</a>	
remote	<a href="#">remote</a>	

#### 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. Optional in the POST or PATCH body
lun_count	integer	The number of LUNs in the application component. Optional in the POST body
name	string	The name of the application component. Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

#### 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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body

Name	Type	Description
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

san

A generic SAN application.

Name	Type	Description
application_components	array[ <a href="#">application_components</a> ]	
new_igroups	array[ <a href="#">san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the db. Optional in the POST or PATCH body

db

Name	Type	Description
size	integer	The size of the db. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the log db. Optional in the POST or PATCH body

log

Name	Type	Description
size	integer	The size of the log db. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

storage\_service

Name	Type	Description
name	string	The storage service of the temp db. Optional in the POST or PATCH body

temp\_db

Name	Type	Description
size	integer	The size of the temp db. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

Name	Type	Description
storage_service	<a href="#">storage_service</a>	

sql\_on\_san

Microsoft SQL using SAN.

Name	Type	Description
db	<a href="#">db</a>	
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. Required in the POST body and optional in the PATCH body
log	<a href="#">log</a>	
new_igroups	array[ <a href="#">sql_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Optional in the POST body
protection_type	<a href="#">protection_type</a>	
server_cores_count	integer	The number of server cores for the db. Optional in the POST body
temp_db	<a href="#">temp_db</a>	

access

Name	Type	Description
installer	string	SQL installer admin user name. Optional in the POST body
service_account	string	SQL service account user name. Required in the POST body

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. Optional in the POST body
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 snapshots
used	integer	The amount of spacing currently in use by the system to store snapshots

## 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"> <li>• readOnly: 1</li> </ul>
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 snapshots 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 snapshots</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 snapshots</p>

Name	Type	Description
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	<a href="#">iops</a>	
latency	<a href="#">latency</a>	
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	<a href="#">snapshot</a>	
space	<a href="#">space</a>	
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	<a href="#">storage_service</a>	
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	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 <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>
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

Name	Type	Description
reserved_unused	integer	The amount of space reserved for system features such as snapshots 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 snapshots
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 snapshots
used_percent	integer	The percentage of the originally provisioned space that is currently being used by the application

#### statistics

Name	Type	Description
components	array[ <a href="#">components</a> ]	
iops	<a href="#">iops</a>	
latency	<a href="#">latency</a>	
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	<a href="#">snapshot</a>	
space	<a href="#">space</a>	

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

#### self\_link

Name	Type	Description
self	<a href="#">href</a>	

#### template

Name	Type	Description
_links	<a href="#">self_link</a>	
name	string	The name of the template that was used to provision this application. Optional in the POST body
protocol	string	The protocol access of the template that was used to provision this application

Name	Type	Description
version	integer	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 <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

#### storage\_service

Name	Type	Description
name	string	The storage service of the desktops. Optional in the POST or PATCH body

#### desktops

Name	Type	Description
count	integer	The number of desktops to support. Optional in the POST or PATCH body
size	integer	The size of the desktops. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body
storage_service	<a href="#">storage_service</a>	

#### hyper\_v\_access

Name	Type	Description
service_account	string	Hyper-V service account. Optional in the POST body

#### vdi\_on\_nas

A VDI application using NAS.

Name	Type	Description
desktops	<a href="#">desktops</a>	
hyper_v_access	<a href="#">hyper_v_access</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

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. Required in the POST body and optional in the PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

vdi\_on\_san

A VDI application using SAN.

Name	Type	Description
desktops	<a href="#">desktops</a>	
hypervisor	string	The name of the hypervisor hosting the application. Required in the POST body
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. Required in the POST body and optional in the PATCH body

Name	Type	Description
new_igroups	array[vdi_on_san_new_igroups]	The list of initiator groups to create. Optional in the POST or PATCH body
protection_type	protection_type	

#### storage\_service

Name	Type	Description
name	string	The storage service of the datastore. Optional in the POST or PATCH body

#### datastore

Name	Type	Description
count	integer	The number of datastores to support. Optional in the POST or PATCH body
size	integer	The size of the datastore. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body
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. Optional in the POST body
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. Required in the POST body and optional in the PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

vsi\_on\_san

A VSI application using SAN.

Name	Type	Description
datastore	<a href="#">datastore</a>	
hypervisor	string	The name of the hypervisor hosting the application. Required in the POST body
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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">vsi_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
protection_type	<a href="#">protection_type</a>	

application

Applications

Name	Type	Description
_links	<a href="#">_links</a>	
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 snapshots. Support for snapshots was added at generation 2. Any future generation numbers and their feature set will be documented
maxdata_on_san	<a href="#">maxdata_on_san</a>	MAX Data application using SAN.
mongo_db_on_san	<a href="#">mongo_db_on_san</a>	MongoDB using SAN.
name	string	Application Name. This field is user supplied when the application is created. Required in the POST body
nas	<a href="#">nas</a>	A generic NAS application.
oracle_on_nfs	<a href="#">oracle_on_nfs</a>	Oracle using NFS.
oracle_on_san	<a href="#">oracle_on_san</a>	Oracle using SAN.
oracle_rac_on_nfs	<a href="#">oracle_rac_on_nfs</a>	Oracle RAC using NFS.
oracle_rac_on_san	<a href="#">oracle_rac_on_san</a>	Oracle RAC using SAN.
protection_granularity	string	Protection granularity determines the scope of Snapshot operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot operations are performed on the entire application. If the value is "component", Snapshot operations are performed separately on the application components
rpo	<a href="#">rpo</a>	
san	<a href="#">san</a>	A generic SAN application.



Name	Type	Description
sql_on_san	<a href="#">sql_on_san</a>	Microsoft SQL using SAN.
sql_on_smb	<a href="#">sql_on_smb</a>	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	<a href="#">statistics</a>	
svm	<a href="#">svm</a>	
template	<a href="#">template</a>	
uuid	string	Application UUID. This field is generated when the application is created. Required in the URL
vdi_on_nas	<a href="#">vdi_on_nas</a>	A VDI application using NAS.
vdi_on_san	<a href="#">vdi_on_san</a>	A VDI application using SAN.
vsi_on_nas	<a href="#">vsi_on_nas</a>	A VSI application using NAS.
vsi_on_san	<a href="#">vsi_on_san</a>	A VSI application using SAN.

#### \_links

Name	Type	Description
self	<a href="#">href</a>	

#### job\_link

Name	Type	Description
<u>_links</u>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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

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.
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.

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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">application_component</a> ]	

## 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": 0,
        "path": "string",
        "size": 0,
        "uuid": "string"
      },
      "volumes": {
        "creaton_timestamp": 0,
        "name": "string",
        "size": 0,
        "uuid": "string"
      }
    },
    "cifs_access": {
      "backing_storage": {
        "type": "volume",
        "uuid": "string"
      },
      "ips": {
```

```

    },
    "path": "string",
    "permissions": {
      "access": "string",
      "user_or_group": "string"
    },
    "server": {
      "name": "string"
    },
    "share": {
      "name": "string"
    }
  },
  "file_system": "mlfs",
  "host_management_url": "string",
  "host_name": "string",
  "name": "string",
  "nfs_access": {
    "backing_storage": {
      "type": "volume",
      "uuid": "string"
    },
    "export_policy": {
      "name": "string"
    },
    "ips": {
    },
    "path": "string",
    "permissions": {
      "access": "string",
      "host": "string"
    }
  },
  "protection_groups": {
    "name": "string",
    "rpo": {
      "local": {
        "description": "string",
        "name": "none"
      },
      "remote": {
        "description": "string",
        "name": "none"
      }
    }
  },
  "uuid": "string"

```

```

},
"san_access": {
  "backing_storage": {
    "type": "volume",
    "uuid": "string"
  },
  "lun_mappings": {
    "fcp": {
      "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": {
      },
      "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

application

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Application name
uuid	string	The application UUID. Valid in URL.

application\_lun\_object

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

application\_volume\_object

Name	Type	Description
creaton_timestamp	date	Creation Time

Name	Type	Description
name	string	Name
size	integer	Size
uuid	string	UUID

#### application\_backing\_storage

Name	Type	Description
luns	array[ <a href="#">application_lun_object</a> ]	
volumes	array[ <a href="#">application_volume_object</a> ]	

#### 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	<a href="#">backing_storage</a>	
ips	array[string]	
path	string	Junction path
permissions	array[ <a href="#">permissions</a> ]	
server	<a href="#">server</a>	
share	<a href="#">share</a>	

#### 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	<a href="#">backing_storage</a>	
export_policy	<a href="#">export_policy</a>	
ips	array[string]	
path	string	Junction path
permissions	array[ <a href="#">permissions</a> ]	

#### local

Name	Type	Description
description	string	A detailed description of the local RPO. This includes details on 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	<a href="#">local</a>	
remote	<a href="#">remote</a>	

application\_protection\_groups

Name	Type	Description
name	string	Protection group name
rpo	<a href="#">rpo</a>	
uuid	string	Protection group UUID

fc\_interface\_reference

An FC interface.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FC interface.
uuid	string	The unique identifier of the FC interface.

Name	Type	Description
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	<a href="#">fc_interface_reference</a>	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

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

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
<a href="#">ip</a>	<a href="#">ip</a>	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	<a href="#">_links</a>	
interface	<a href="#">interface</a>	A network interface. Either UUID or name may be supplied on input.
port	integer	The TCP port number of the iSCSI access endpoint.

#### application\_lun\_mapping\_object

Name	Type	Description
fc	array[ <a href="#">application_san_access_fcp_endpoint</a> ]	All possible Fibre Channel Protocol (FCP) access endpoints for the LUN.
igroup	<a href="#">igroup</a>	
iscsi	array[ <a href="#">application_san_access_iscsi_endpoint</a> ]	All possible iSCSI access endpoints for the LUN.
lun_id	integer	LUN ID

#### application\_san\_access

Name	Type	Description
backing_storage	<a href="#">backing_storage</a>	
is_clone	boolean	Clone
lun_mappings	array[ <a href="#">application_lun_mapping_object</a> ]	
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

Name	Type	Description
_links	<a href="#">_links</a>	
application	<a href="#">application</a>	
backing_storage	<a href="#">application_backing_storage</a>	
cifs_access	array[ <a href="#">application_cifs_properties</a> ]	
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[ <a href="#">application_nfs_properties</a> ]	
protection_groups	array[ <a href="#">application_protection_groups</a> ]	
san_access	array[ <a href="#">application_san_access</a> ]	
storage_service	<a href="#">storage_service</a>	
svm	<a href="#">svm</a>	
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[ <a href="#">error_arguments</a> ]	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

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



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

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
num_records	integer	Number of records
records	array[ <a href="#">application_component_snapshot</a> ]	

## 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": "crash",
    "create_time": "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

application

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Application Name
uuid	string	Application UUID. Valid in URL

component

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

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create an application component Snapshot copy

POST

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

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

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
application	<a href="#">application</a>	
comment	string	Comment. Valid in POST
component	<a href="#">component</a>	

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



## Example request

```
{
  "_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": "crash",
  "create_time": "string",
  "svm": {
    "name": "string",
    "uuid": "string"
  },
  "uuid": "string"
}
```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

application

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Application Name
uuid	string	Application UUID. Valid in URL

component

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
application	<a href="#">application</a>	
comment	string	Comment. Valid in POST

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

#### job\_link

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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 component Snapshot copy

DELETE

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

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	
component.uuid	string	path	True	
uuid	string	path	True	

### Response

```
Status: 202, Accepted
```

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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	<a href="#">_links</a>	
application	<a href="#">application</a>	
comment	string	Comment. Valid in POST
component	<a href="#">component</a>	
consistency_type	string	Consistency Type. This is for categorization only. A snapshot should not be set to application consistent unless the host application is quiesced for the snapshot. Valid in POST



Name	Type	Description
create_time	string	Creation Time
is_partial	boolean	A partial snapshot means that not all volumes in an application component were included in the snapshot.
name	string	Snapshot Name. Valid in POST
svm	<a href="#">svm</a>	
uuid	string	Snapshot 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": "crash",
  "create_time": "string",
  "svm": {
    "name": "string",
    "uuid": "string"
  },
  "uuid": "string"
}
```

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

application

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Application Name
uuid	string	Application UUID. Valid in URL

component

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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`

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

### Parameters

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

### 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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
<code>_links</code>	<a href="#">_links</a>	
<code>application</code>	<a href="#">application</a>	
<code>backing_storage</code>	<a href="#">application_backing_storage</a>	
<code>cifs_access</code>	array[ <a href="#">application_cifs_properties</a> ]	
<code>file_system</code>	string	Defines the type of file system that will be installed on this application component.



<b>Name</b>	<b>Type</b>	<b>Description</b>
host_management_url	string	Host management URL
host_name	string	L2 Host FQDN
name	string	Application component name
nfs_access	array[ <a href="#">application_nfs_properties</a> ]	
protection_groups	array[ <a href="#">application_protection_groups</a> ]	
san_access	array[ <a href="#">application_san_access</a> ]	
storage_service	<a href="#">storage_service</a>	
svm	<a href="#">svm</a>	
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": 0,
      "path": "string",
      "size": 0,
      "uuid": "string"
    },
    "volumes": {
      "creation_timestamp": 0,
      "name": "string",
      "size": 0,
      "uuid": "string"
    }
  },
  "cifs_access": {
    "backing_storage": {
      "type": "volume",
      "uuid": "string"
    },
    "ips": {
    },
    "path": "string",
    "permissions": {
      "access": "string",
      "user_or_group": "string"
    },
    "server": {
      "name": "string"
    }
  },
}
```

```

    "share": {
      "name": "string"
    }
  },
  "file_system": "mlfs",
  "host_management_url": "string",
  "host_name": "string",
  "name": "string",
  "nfs_access": {
    "backing_storage": {
      "type": "volume",
      "uuid": "string"
    },
    "export_policy": {
      "name": "string"
    },
    "ips": {
    },
    "path": "string",
    "permissions": {
      "access": "string",
      "host": "string"
    }
  },
  "protection_groups": {
    "name": "string",
    "rpo": {
      "local": {
        "description": "string",
        "name": "none"
      },
      "remote": {
        "description": "string",
        "name": "none"
      }
    },
    "uuid": "string"
  },
  "san_access": {
    "backing_storage": {
      "type": "volume",
      "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": {
  },
  "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	<a href="#">href</a>	

application

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Application name
uuid	string	The application UUID. Valid in URL.

application\_lun\_object

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

application\_volume\_object

Name	Type	Description
creaton_timestamp	date	Creation Time
name	string	Name
size	integer	Size
uuid	string	UUID

### application\_backing\_storage

Name	Type	Description
luns	array[ <a href="#">application_lun_object</a> ]	
volumes	array[ <a href="#">application_volume_object</a> ]	

### 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	<a href="#">backing_storage</a>	
ips	array[string]	
path	string	Junction path
permissions	array[ <a href="#">permissions</a> ]	
server	<a href="#">server</a>	
share	<a href="#">share</a>	

## 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	<a href="#">backing_storage</a>	
export_policy	<a href="#">export_policy</a>	
ips	array[string]	
path	string	Junction path
permissions	array[ <a href="#">permissions</a> ]	

## 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	Type	Description
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	<a href="#">local</a>	
remote	<a href="#">remote</a>	

application\_protection\_groups

Name	Type	Description
name	string	Protection group name
rpo	<a href="#">rpo</a>	
uuid	string	Protection group UUID

fc\_interface\_reference

An FC interface.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">fc_interface_reference</a>	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

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

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
ip	<a href="#">ip</a>	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
<a href="#">_links</a>	<a href="#">_links</a>	
interface	<a href="#">interface</a>	A network interface. Either UUID or name may be supplied on input.
port	integer	The TCP port number of the iSCSI access endpoint.

application\_lun\_mapping\_object

Name	Type	Description
fc	array[ <a href="#">application_san_access_fc_endpoint</a> ]	All possible Fibre Channel Protocol (FCP) access endpoints for the LUN.
igroup	<a href="#">igroup</a>	
iscsi	array[ <a href="#">application_san_access_iscsi_endpoint</a> ]	All possible iSCSI access endpoints for the LUN.
lun_id	integer	LUN ID

#### application\_san\_access

Name	Type	Description
backing_storage	<a href="#">backing_storage</a>	
is_clone	boolean	Clone
lun_mappings	array[ <a href="#">application_lun_mapping_object</a> ]	
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[ <a href="#">error_arguments</a> ]	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
```

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..2017-05-04T23:59:59-05:00
```

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

GET

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

#### 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
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.

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

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">application_snapshot</a> ]	

## 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": "crash",
    "create_time": "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

application

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Application name
uuid	string	The application UUID. Valid in URL.

components

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
application	<a href="#">application</a>	
comment	string	Comment. Valid in POST.
components	array[ <a href="#">components</a> ]	
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	<a href="#">svm</a>	
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[ <a href="#">error_arguments</a> ]	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

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

#### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
application	<a href="#">application</a>	
comment	string	Comment. Valid in POST.
components	array[ <a href="#">components</a> ]	
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	<a href="#">svm</a>	
uuid	string	The Snapshot copy UUID. Valid in URL.

## Example request

```
{
  "_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": "crash",
  "create_time": "string",
  "svm": {
    "name": "string",
    "uuid": "string"
  },
  "uuid": "string"
}
```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

application

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Application name
uuid	string	The application UUID. Valid in URL.

components

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
application	<a href="#">application</a>	

Name	Type	Description
comment	string	Comment. Valid in POST.
components	array[ <a href="#">components</a> ]	
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	<a href="#">svm</a>	
uuid	string	The Snapshot copy UUID. Valid in URL.

#### job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
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}

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
```

#### Learn more

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

#### Parameters

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

#### Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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
<a href="#">_links</a>	<a href="#">_links</a>	
<a href="#">application</a>	<a href="#">application</a>	
comment	string	Comment. Valid in POST.
components	array[ <a href="#">components</a> ]	
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	<a href="#">svm</a>	
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": "crash",
  "create_time": "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	<a href="#">href</a>	

application

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Application name
uuid	string	The application UUID. Valid in URL.

components

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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

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

#### Parameters

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

#### Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

### Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

### 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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 [DOC 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
<a href="#">_links</a>	<a href="#">_links</a>	
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 snapshots. Support for snapshots was added at generation 2. Any future generation numbers and their feature set will be documented
maxdata_on_san	<a href="#">maxdata_on_san</a>	MAX Data application using SAN.
mongo_db_on_san	<a href="#">mongo_db_on_san</a>	MongoDB using SAN.
name	string	Application Name. This field is user supplied when the application is created. Required in the POST body
nas	<a href="#">nas</a>	A generic NAS application.
oracle_on_nfs	<a href="#">oracle_on_nfs</a>	Oracle using NFS.
oracle_on_san	<a href="#">oracle_on_san</a>	Oracle using SAN.
oracle_rac_on_nfs	<a href="#">oracle_rac_on_nfs</a>	Oracle RAC using NFS.
oracle_rac_on_san	<a href="#">oracle_rac_on_san</a>	Oracle RAC using SAN.



Name	Type	Description
protection_granularity	string	Protection granularity determines the scope of Snapshot operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot operations are performed on the entire application. If the value is "component", Snapshot operations are performed separately on the application components
rpo	<a href="#">rpo</a>	
san	<a href="#">san</a>	A generic SAN application.
sql_on_san	<a href="#">sql_on_san</a>	Microsoft SQL using SAN.
sql_on_smb	<a href="#">sql_on_smb</a>	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	<a href="#">statistics</a>	
svm	<a href="#">svm</a>	
template	<a href="#">template</a>	
uuid	string	Application UUID. This field is generated when the application is created. Required in the URL
vdi_on_nas	<a href="#">vdi_on_nas</a>	A VDI application using NAS.
vdi_on_san	<a href="#">vdi_on_san</a>	A VDI application using SAN.
vsi_on_nas	<a href="#">vsi_on_nas</a>	A VSI application using NAS.
vsi_on_san	<a href="#">vsi_on_san</a>	A VSI application using SAN.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    },
    "snapshots": {
      "href": "/api/resourcelink"
    }
  },
  "creation_timestamp": "string",
  "generation": 0,
  "maxdata_on_san": {
    "app_type": "mongodb",
    "application_components": {
      "file_system": "generic",
      "host_management_url": "string",
      "metadata": {
      },
      "protection_type": {
        "local_rpo": "6_hourly",
        "remote_rpo": "6_hourly"
      },
      "storage_service": {
        "name": "extreme"
      }
    },
    "metadata": {
    },
    "new_igroups": {
      "initiators": {
      },
      "os_type": "aix",
      "protocol": "fc"
    },
    "ocsm_url": "string",
    "os_type": "aix"
  },
  "mongo_db_on_san": {
    "dataset": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "new_igroups": {
```

```

    "initiators": {
    },
    "os_type": "hyper_v",
    "protocol": "fcp"
  },
  "os_type": "hyper_v",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "secondary_igroups": {
  }
},
"nas": {
  "application_components": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "cifs_access": {
    "access": "change"
  },
  "nfs_access": {
    "access": "none"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"oracle_on_nfs": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "ora_home": {
    "storage_service": {

```

```

    "name": "extreme"
  }
},
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
},
"redo_log": {
  "storage_service": {
    "name": "extreme"
  }
}
},
"oracle_on_san": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fc"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
}
},

```

```

"oracle_rac_on_nfs": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "grid_binary": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_crs": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_rac_on_san": {
    "archive_log": {
      "storage_service": {
        "name": "extreme"
      }
    }
  },
  "db": {

```

```

    "storage_service": {
      "name": "extreme"
    }
  },
  "db_sids": {
  },
  "grid_binary": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fcp"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_crs": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"protection_granularity": "application",
"rpo": {
  "components": {
    "name": "string",
    "rpo": {
      "local": {
        "description": "string",
        "name": "6_hourly"
      }
    }
  }
}

```

```

    },
    "remote": {
      "description": "string",
      "name": "6_hourly"
    }
  },
  "uuid": "string"
},
"local": {
  "description": "string",
  "name": "6_hourly"
},
"remote": {
  "description": "string",
  "name": "6_hourly"
}
},
"san": {
  "application_components": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fcp"
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"sql_on_san": {
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "log": {
    "storage_service": {
      "name": "extreme"
    }
  }
},

```

```

"new_igroups": {
  "initiators": {
  },
  "os_type": "hyper_v",
  "protocol": "fc"
},
"os_type": "windows",
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
},
"temp_db": {
  "storage_service": {
    "name": "extreme"
  }
}
},
"sql_on_smb": {
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "temp_db": {
    "storage_service": {
      "name": "extreme"
    }
  }
}
},
"state": "creating",
"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
}

```

```

},
"template": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"protocol": "nas",
"version": 0
},
"uuid": "string",
"vdi_on_nas": {
  "desktops": {
    "storage_service": {
      "name": "extreme"
    }
  },
},
"nfs_access": {
  "access": "none"
},
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
}
},
"vdi_on_san": {
  "desktops": {
    "storage_service": {
      "name": "extreme"
    }
  },
},
"hypervisor": "hyper_v",
"new_igroups": {
  "initiators": {
  },
  "protocol": "fc"
},
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
}
},
"vsi_on_nas": {
  "datastore": {
    "storage_service": {
      "name": "extreme"
    }
  }
}

```

```

    }
  },
  "nfs_access": {
    "access": "none"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"vsi_on_san": {
  "datastore": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "hypervisor": "hyper_v",
  "new_igroups": {
    "initiators": {
    },
    "protocol": "fcp"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
}
}
}

```

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
snapshots	<a href="#">href</a>	

metadata

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

protection\_type

Name	Type	Description
local_rpo	string	The local rpo of the application component. Optional in the POST or PATCH body
remote_rpo	string	The remote rpo of the application component. Optional in the POST or PATCH body

storage\_service

Name	Type	Description
name	string	The storage service of the application component. Optional in the POST or PATCH body

maxdata\_on\_san\_application\_components

application-components

Name	Type	Description
file_system	string	Defines the kind of file system that will be installed on this application component. Optional in the POST body
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. Required in the POST body
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. Required in the POST body and optional in the PATCH body
lun_count	integer	The number of LUNs in the application component. Required in the POST body
metadata	array[ <a href="#">metadata</a> ]	
name	string	The name of the application component. Required in the POST body and optional in the PATCH body
protection_type	<a href="#">protection_type</a>	
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body

metadata

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

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

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. Required in the POST body
application_components	array[ <a href="#">maxdata_on_san_application_components</a> ]	application-components. Optional in the POST or PATCH body
metadata	array[ <a href="#">metadata</a> ]	
new_igroups	array[ <a href="#">maxdata_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body

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. Required in the POST body

#### storage\_service

Name	Type	Description
name	string	The storage service of the database. Optional in the POST or PATCH body

#### 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. Optional in the POST body
replication_factor	integer	The number of data bearing members of the replicaset, including 1 primary and at least 1 secondary. Optional in the POST body
size	integer	The size of the database. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

#### mongo\_db\_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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

#### protection\_type

Name	Type	Description
local_rpo	string	The local rpo of the application. Optional in the POST or PATCH body
remote_rpo	string	The remote rpo of the application. Optional in the POST body

#### secondary\_igroups

Name	Type	Description
name	string	The name of the initiator group for each secondary. Optional in the POST or PATCH body

#### mongo\_db\_on\_san

MongoDB using SAN.

Name	Type	Description
dataset	<a href="#">dataset</a>	
new_igroups	array[ <a href="#">mongo_db_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body

Name	Type	Description
os_type	string	The name of the host OS running the application. Optional in the POST body
primary_igroup_name	string	The initiator group for the primary. Required in the POST body and optional in the PATCH body
protection_type	<a href="#">protection_type</a>	
secondary_igroups	array[ <a href="#">secondary_igroups</a> ]	

#### application\_components

Name	Type	Description
name	string	The name of the application component. Optional in the POST or PATCH body
share_count	integer	The number of shares in the application component. Optional in the POST body
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member shares. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

#### app\_cifs\_access

The list of CIFS access controls.

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

#### app\_nfs\_access

The list of NFS access controls.

Name	Type	Description
access	string	The NFS access granted. Optional in the POST body
host	string	The name of the NFS entity granted access. Optional in the POST body

nas

A generic NAS application.

Name	Type	Description
application_components	array[ <a href="#">application_components</a> ]	
cifs_access	array[ <a href="#">app_cifs_access</a> ]	The list of CIFS access controls. Optional in the POST body
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the archive log. Optional in the POST or PATCH body

archive\_log

Name	Type	Description
size	integer	The size of the archive log. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

db

Name	Type	Description
size	integer	The size of the database. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the ORACLE_HOME storage volume. Optional in the POST or PATCH body

ora\_home

Name	Type	Description
size	integer	The size of the ORACLE_HOME storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the redo log group. Optional in the POST or PATCH body

redo\_log

Name	Type	Description
mirrored	boolean	Should the redo log group be mirrored? Optional in the POST body

Name	Type	Description
size	integer	The size of the redo log group. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

oracle\_on\_nfs

Oracle using NFS.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
ora_home	<a href="#">ora_home</a>	
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

oracle\_on\_san

Oracle using SAN.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">oracle_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
ora_home	<a href="#">ora_home</a>	
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the Oracle grid binary storage volume. Optional in the POST or PATCH body

grid\_binary

Name	Type	Description
size	integer	The size of the Oracle grid binary storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

Name	Type	Description
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the Oracle CRS volume. Optional in the POST or PATCH body

oracle\_crs

Name	Type	Description
copies	integer	The number of CRS volumes. Optional in the POST body
size	integer	The size of the Oracle CRS/voting storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST body
storage_service	<a href="#">storage_service</a>	

oracle\_rac\_on\_nfs

Oracle RAC using NFS.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
grid_binary	<a href="#">grid_binary</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
ora_home	<a href="#">ora_home</a>	
oracle_crs	<a href="#">oracle_crs</a>	
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

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. Optional in the POST or PATCH body

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

oracle\_rac\_on\_san

Oracle RAC using SAN.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
db_sids	array[ <a href="#">db_sids</a> ]	
grid_binary	<a href="#">grid_binary</a>	



Name	Type	Description
new_igroups	array[ <a href="#">oracle_rac_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
ora_home	<a href="#">ora_home</a>	
oracle_crs	<a href="#">oracle_crs</a>	
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

#### local

Name	Type	Description
description	string	A detailed description of the local RPO. This will include details about the snapshot schedule
name	string	The local RPO of the component. This indicates how often component snapshots 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	<a href="#">local</a>	
remote	<a href="#">remote</a>	

#### components

Name	Type	Description
name	string	Component Name
rpo	<a href="#">rpo</a>	
uuid	string	Component UUID

#### local

Name	Type	Description
description	string	A detailed description of the local RPO. This will include details about the snapshot schedule
name	string	The local RPO of the application. This indicates how often application snapshots 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[ <a href="#">components</a> ]	
is_supported	boolean	Is RPO supported for this application? Generation 1 applications did not support snapshots or MetroCluster
local	<a href="#">local</a>	
remote	<a href="#">remote</a>	

#### 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. Optional in the POST or PATCH body
lun_count	integer	The number of LUNs in the application component. Optional in the POST body
name	string	The name of the application component. Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

#### 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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body

Name	Type	Description
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

san

A generic SAN application.

Name	Type	Description
application_components	array[ <a href="#">application_components</a> ]	
new_igroups	array[ <a href="#">san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the db. Optional in the POST or PATCH body

db

Name	Type	Description
size	integer	The size of the db. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the log db. Optional in the POST or PATCH body

log

Name	Type	Description
size	integer	The size of the log db. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

storage\_service

Name	Type	Description
name	string	The storage service of the temp db. Optional in the POST or PATCH body

temp\_db

Name	Type	Description
size	integer	The size of the temp db. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

Name	Type	Description
storage_service	<a href="#">storage_service</a>	

sql\_on\_san

Microsoft SQL using SAN.

Name	Type	Description
db	<a href="#">db</a>	
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. Required in the POST body and optional in the PATCH body
log	<a href="#">log</a>	
new_igroups	array[ <a href="#">sql_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Optional in the POST body
protection_type	<a href="#">protection_type</a>	
server_cores_count	integer	The number of server cores for the db. Optional in the POST body
temp_db	<a href="#">temp_db</a>	

access

Name	Type	Description
installer	string	SQL installer admin user name. Optional in the POST body
service_account	string	SQL service account user name. Required in the POST body

sql\_on\_smb

## Microsoft SQL using SMB.

Name	Type	Description
access	<a href="#">access</a>	
db	<a href="#">db</a>	
log	<a href="#">log</a>	
protection_type	<a href="#">protection_type</a>	
server_cores_count	integer	The number of server cores for the db. Optional in the POST body
temp_db	<a href="#">temp_db</a>	

## 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 snapshots
used	integer	The amount of spacing currently in use by the system to store snapshots

## 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"> <li>• readOnly: 1</li> </ul>
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 snapshots 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 snapshots</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 snapshots</p>



Name	Type	Description
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	<a href="#">iops</a>	
latency	<a href="#">latency</a>	
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	<a href="#">snapshot</a>	
space	<a href="#">space</a>	
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	<a href="#">storage_service</a>	
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	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 <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>
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

Name	Type	Description
reserved_unused	integer	The amount of space reserved for system features such as snapshots 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 snapshots
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 snapshots
used_percent	integer	The percentage of the originally provisioned space that is currently being used by the application

#### statistics

Name	Type	Description
components	array[ <a href="#">components</a> ]	
iops	<a href="#">iops</a>	
latency	<a href="#">latency</a>	
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	<a href="#">snapshot</a>	
space	<a href="#">space</a>	

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

self\_link

Name	Type	Description
self	<a href="#">href</a>	

template

Name	Type	Description
_links	<a href="#">self_link</a>	
name	string	The name of the template that was used to provision this application. Optional in the POST body
protocol	string	The protocol access of the template that was used to provision this application

Name	Type	Description
version	integer	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 <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

#### storage\_service

Name	Type	Description
name	string	The storage service of the desktops. Optional in the POST or PATCH body

#### desktops

Name	Type	Description
count	integer	The number of desktops to support. Optional in the POST or PATCH body
size	integer	The size of the desktops. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body
storage_service	<a href="#">storage_service</a>	

#### hyper\_v\_access

Name	Type	Description
service_account	string	Hyper-V service account. Optional in the POST body

#### vdi\_on\_nas

A VDI application using NAS.

Name	Type	Description
desktops	<a href="#">desktops</a>	
hyper_v_access	<a href="#">hyper_v_access</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

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. Required in the POST body and optional in the PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

vdi\_on\_san

A VDI application using SAN.

Name	Type	Description
desktops	<a href="#">desktops</a>	
hypervisor	string	The name of the hypervisor hosting the application. Required in the POST body
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. Required in the POST body and optional in the PATCH body

Name	Type	Description
new_igroups	array[ <a href="#">vdi_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
protection_type	<a href="#">protection_type</a>	

#### storage\_service

Name	Type	Description
name	string	The storage service of the datastore. Optional in the POST or PATCH body

#### datastore

Name	Type	Description
count	integer	The number of datastores to support. Optional in the POST or PATCH body
size	integer	The size of the datastore. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body
storage_service	<a href="#">storage_service</a>	

#### vsi\_on\_nas

A VSI application using NAS.

Name	Type	Description
datastore	<a href="#">datastore</a>	
hyper_v_access	<a href="#">hyper_v_access</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

#### 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. Required in the POST body and optional in the PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

vsi\_on\_san

A VSI application using SAN.

Name	Type	Description
datastore	<a href="#">datastore</a>	
hypervisor	string	The name of the hypervisor hosting the application. Required in the POST body
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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">vsi_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
protection_type	<a href="#">protection_type</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error



Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}

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

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Application UUID

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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 snapshots. Support for snapshots was added at generation 2. Any future generation numbers and their feature set will be documented
maxdata_on_san	<a href="#">maxdata_on_san</a>	MAX Data application using SAN.
mongo_db_on_san	<a href="#">mongo_db_on_san</a>	MongoDB using SAN.
name	string	Application Name. This field is user supplied when the application is created. Required in the POST body
nas	<a href="#">nas</a>	A generic NAS application.
oracle_on_nfs	<a href="#">oracle_on_nfs</a>	Oracle using NFS.
oracle_on_san	<a href="#">oracle_on_san</a>	Oracle using SAN.
oracle_rac_on_nfs	<a href="#">oracle_rac_on_nfs</a>	Oracle RAC using NFS.
oracle_rac_on_san	<a href="#">oracle_rac_on_san</a>	Oracle RAC using SAN.

Name	Type	Description
protection_granularity	string	Protection granularity determines the scope of Snapshot operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot operations are performed on the entire application. If the value is "component", Snapshot operations are performed separately on the application components
rpo	rpo	
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.
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. Required in the URL
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 request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    },
    "snapshots": {
      "href": "/api/resourcelink"
    }
  },
  "creation_timestamp": "string",
  "generation": 0,
  "maxdata_on_san": {
    "app_type": "mongodb",
    "application_components": {
      "file_system": "generic",
      "host_management_url": "string",
      "metadata": {
      },
      "protection_type": {
        "local_rpo": "6_hourly",
        "remote_rpo": "6_hourly"
      },
      "storage_service": {
        "name": "extreme"
      }
    },
    "metadata": {
    },
    "new_igroups": {
      "initiators": {
      },
      "os_type": "aix",
      "protocol": "fc"
    },
    "ocsm_url": "string",
    "os_type": "aix"
  },
  "mongo_db_on_san": {
    "dataset": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "new_igroups": {
```

```

    "initiators": {
    },
    "os_type": "hyper_v",
    "protocol": "fcp"
  },
  "os_type": "hyper_v",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "secondary_igroups": {
  }
},
"nas": {
  "application_components": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "cifs_access": {
    "access": "change"
  },
  "nfs_access": {
    "access": "none"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"oracle_on_nfs": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "ora_home": {
    "storage_service": {

```

```

    "name": "extreme"
  }
},
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
},
"redo_log": {
  "storage_service": {
    "name": "extreme"
  }
}
},
"oracle_on_san": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fc"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
}
},

```

```

"oracle_rac_on_nfs": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "grid_binary": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_crs": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_rac_on_san": {
    "archive_log": {
      "storage_service": {
        "name": "extreme"
      }
    }
  },
  "db": {

```



```

    "storage_service": {
      "name": "extreme"
    }
  },
  "db_sids": {
  },
  "grid_binary": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fcp"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_crs": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"protection_granularity": "application",
"rpo": {
  "components": {
    "name": "string",
    "rpo": {
      "local": {
        "description": "string",
        "name": "6_hourly"
      }
    }
  }
}

```

```

    },
    "remote": {
      "description": "string",
      "name": "6_hourly"
    }
  },
  "uuid": "string"
},
"local": {
  "description": "string",
  "name": "6_hourly"
},
"remote": {
  "description": "string",
  "name": "6_hourly"
}
},
"san": {
  "application_components": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fc"
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"sql_on_san": {
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "log": {
    "storage_service": {
      "name": "extreme"
    }
  }
},

```

```

"new_igroups": {
  "initiators": {
  },
  "os_type": "hyper_v",
  "protocol": "fcp"
},
"os_type": "windows",
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
},
"temp_db": {
  "storage_service": {
    "name": "extreme"
  }
}
},
"sql_on_smb": {
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "temp_db": {
    "storage_service": {
      "name": "extreme"
    }
  }
}
},
"state": "creating",
"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
}

```

```

},
"template": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"protocol": "nas",
"version": 0
},
"uuid": "string",
"vdi_on_nas": {
  "desktops": {
    "storage_service": {
      "name": "extreme"
    }
  },
},
"nfs_access": {
  "access": "none"
},
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
}
},
"vdi_on_san": {
  "desktops": {
    "storage_service": {
      "name": "extreme"
    }
  },
},
"hypervisor": "hyper_v",
"new_igroups": {
  "initiators": {
  },
  "protocol": "fcp"
},
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
}
},
"vsi_on_nas": {
  "datastore": {
    "storage_service": {
      "name": "extreme"
    }
  }
}

```

```

    }
  },
  "nfs_access": {
    "access": "none"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"vsi_on_san": {
  "datastore": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "hypervisor": "hyper_v",
  "new_igroups": {
    "initiators": {
    },
    "protocol": "fc"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
}
}
}

```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	
snapshots	<a href="#">href</a>	

metadata

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

protection\_type

Name	Type	Description
local_rpo	string	The local rpo of the application component. Optional in the POST or PATCH body
remote_rpo	string	The remote rpo of the application component. Optional in the POST or PATCH body

storage\_service

Name	Type	Description
name	string	The storage service of the application component. Optional in the POST or PATCH body

maxdata\_on\_san\_application\_components

application-components



Name	Type	Description
file_system	string	Defines the kind of file system that will be installed on this application component. Optional in the POST body
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. Required in the POST body
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. Required in the POST body and optional in the PATCH body
lun_count	integer	The number of LUNs in the application component. Required in the POST body
metadata	array[ <a href="#">metadata</a> ]	
name	string	The name of the application component. Required in the POST body and optional in the PATCH body
protection_type	<a href="#">protection_type</a>	
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body

metadata

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

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

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. Required in the POST body
application_components	array[ <a href="#">maxdata_on_san_application_components</a> ]	application-components. Optional in the POST or PATCH body
metadata	array[ <a href="#">metadata</a> ]	
new_igroups	array[ <a href="#">maxdata_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body

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. Required in the POST body

#### storage\_service

Name	Type	Description
name	string	The storage service of the database. Optional in the POST or PATCH body

#### 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. Optional in the POST body
replication_factor	integer	The number of data bearing members of the replicaset, including 1 primary and at least 1 secondary. Optional in the POST body
size	integer	The size of the database. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

#### mongo\_db\_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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

#### protection\_type

Name	Type	Description
local_rpo	string	The local rpo of the application. Optional in the POST or PATCH body
remote_rpo	string	The remote rpo of the application. Optional in the POST body

#### secondary\_igroups

Name	Type	Description
name	string	The name of the initiator group for each secondary. Optional in the POST or PATCH body

#### mongo\_db\_on\_san

MongoDB using SAN.

Name	Type	Description
dataset	<a href="#">dataset</a>	
new_igroups	array[ <a href="#">mongo_db_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body

Name	Type	Description
os_type	string	The name of the host OS running the application. Optional in the POST body
primary_igroup_name	string	The initiator group for the primary. Required in the POST body and optional in the PATCH body
protection_type	<a href="#">protection_type</a>	
secondary_igroups	array[ <a href="#">secondary_igroups</a> ]	

#### application\_components

Name	Type	Description
name	string	The name of the application component. Optional in the POST or PATCH body
share_count	integer	The number of shares in the application component. Optional in the POST body
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member shares. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

#### app\_cifs\_access

The list of CIFS access controls.

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

#### app\_nfs\_access

The list of NFS access controls.

Name	Type	Description
access	string	The NFS access granted. Optional in the POST body
host	string	The name of the NFS entity granted access. Optional in the POST body

nas

A generic NAS application.

Name	Type	Description
application_components	array[ <a href="#">application_components</a> ]	
cifs_access	array[ <a href="#">app_cifs_access</a> ]	The list of CIFS access controls. Optional in the POST body
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the archive log. Optional in the POST or PATCH body

archive\_log

Name	Type	Description
size	integer	The size of the archive log. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

db

Name	Type	Description
size	integer	The size of the database. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the ORACLE_HOME storage volume. Optional in the POST or PATCH body

ora\_home

Name	Type	Description
size	integer	The size of the ORACLE_HOME storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the redo log group. Optional in the POST or PATCH body

redo\_log

Name	Type	Description
mirrored	boolean	Should the redo log group be mirrored? Optional in the POST body

Name	Type	Description
size	integer	The size of the redo log group. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

oracle\_on\_nfs

Oracle using NFS.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
ora_home	<a href="#">ora_home</a>	
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

oracle\_on\_san



Oracle using SAN.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">oracle_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
ora_home	<a href="#">ora_home</a>	
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the Oracle grid binary storage volume. Optional in the POST or PATCH body

grid\_binary

Name	Type	Description
size	integer	The size of the Oracle grid binary storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

Name	Type	Description
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the Oracle CRS volume. Optional in the POST or PATCH body

oracle\_crs

Name	Type	Description
copies	integer	The number of CRS volumes. Optional in the POST body
size	integer	The size of the Oracle CRS/voting storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST body
storage_service	<a href="#">storage_service</a>	

oracle\_rac\_on\_nfs

Oracle RAC using NFS.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
grid_binary	<a href="#">grid_binary</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
ora_home	<a href="#">ora_home</a>	
oracle_crs	<a href="#">oracle_crs</a>	
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

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. Optional in the POST or PATCH body

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

oracle\_rac\_on\_san

Oracle RAC using SAN.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
db_sids	array[ <a href="#">db_sids</a> ]	
grid_binary	<a href="#">grid_binary</a>	

Name	Type	Description
new_igroups	array[oracle_rac_on_san_new_igroups]	The list of initiator groups to create. Optional in the POST or PATCH body
ora_home	ora_home	
oracle_crs	oracle_crs	
os_type	string	The name of the host OS running the application. Required in the POST body
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 schedule
name	string	The local RPO of the component. This indicates how often component snapshots 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	<a href="#">rpo</a>	
uuid	string	Component UUID

#### local

Name	Type	Description
description	string	A detailed description of the local RPO. This will include details about the snapshot schedule
name	string	The local RPO of the application. This indicates how often application snapshots 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[ <a href="#">components</a> ]	
is_supported	boolean	Is RPO supported for this application? Generation 1 applications did not support snapshots or MetroCluster
local	<a href="#">local</a>	
remote	<a href="#">remote</a>	

#### 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. Optional in the POST or PATCH body
lun_count	integer	The number of LUNs in the application component. Optional in the POST body
name	string	The name of the application component. Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

#### 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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body

Name	Type	Description
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

san

A generic SAN application.

Name	Type	Description
application_components	array[ <a href="#">application_components</a> ]	
new_igroups	array[ <a href="#">san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the db. Optional in the POST or PATCH body

db

Name	Type	Description
size	integer	The size of the db. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the log db. Optional in the POST or PATCH body

log

Name	Type	Description
size	integer	The size of the log db. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

storage\_service

Name	Type	Description
name	string	The storage service of the temp db. Optional in the POST or PATCH body

temp\_db

Name	Type	Description
size	integer	The size of the temp db. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body



Name	Type	Description
storage_service	<a href="#">storage_service</a>	

sql\_on\_san

Microsoft SQL using SAN.

Name	Type	Description
db	<a href="#">db</a>	
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. Required in the POST body and optional in the PATCH body
log	<a href="#">log</a>	
new_igroups	array[ <a href="#">sql_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Optional in the POST body
protection_type	<a href="#">protection_type</a>	
server_cores_count	integer	The number of server cores for the db. Optional in the POST body
temp_db	<a href="#">temp_db</a>	

access

Name	Type	Description
installer	string	SQL installer admin user name. Optional in the POST body
service_account	string	SQL service account user name. Required in the POST body

sql\_on\_smb

Microsoft SQL using SMB.

Name	Type	Description
access	<a href="#">access</a>	
db	<a href="#">db</a>	
log	<a href="#">log</a>	
protection_type	<a href="#">protection_type</a>	
server_cores_count	integer	The number of server cores for the db. Optional in the POST body
temp_db	<a href="#">temp_db</a>	

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 snapshots
used	integer	The amount of spacing currently in use by the system to store snapshots

## 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"> <li>• readOnly: 1</li> </ul>
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 snapshots 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 snapshots</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 snapshots</p>

Name	Type	Description
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	<a href="#">iops</a>	
latency	<a href="#">latency</a>	
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	<a href="#">snapshot</a>	
space	<a href="#">space</a>	
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	<a href="#">storage_service</a>	
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	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 <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>
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

Name	Type	Description
reserved_unused	integer	The amount of space reserved for system features such as snapshots 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 snapshots
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 snapshots
used_percent	integer	The percentage of the originally provisioned space that is currently being used by the application

#### statistics

Name	Type	Description
components	array[ <a href="#">components</a> ]	
iops	<a href="#">iops</a>	
latency	<a href="#">latency</a>	
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	<a href="#">snapshot</a>	
space	<a href="#">space</a>	

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

self\_link

Name	Type	Description
self	<a href="#">href</a>	

template

Name	Type	Description
_links	<a href="#">self_link</a>	
name	string	The name of the template that was used to provision this application. Optional in the POST body
protocol	string	The protocol access of the template that was used to provision this application

Name	Type	Description
version	integer	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 <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

#### storage\_service

Name	Type	Description
name	string	The storage service of the desktops. Optional in the POST or PATCH body

#### desktops

Name	Type	Description
count	integer	The number of desktops to support. Optional in the POST or PATCH body
size	integer	The size of the desktops. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body
storage_service	<a href="#">storage_service</a>	

#### hyper\_v\_access

Name	Type	Description
service_account	string	Hyper-V service account. Optional in the POST body

#### vdi\_on\_nas

A VDI application using NAS.



Name	Type	Description
desktops	<a href="#">desktops</a>	
hyper_v_access	<a href="#">hyper_v_access</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

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. Required in the POST body and optional in the PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

vdi\_on\_san

A VDI application using SAN.

Name	Type	Description
desktops	<a href="#">desktops</a>	
hypervisor	string	The name of the hypervisor hosting the application. Required in the POST body
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. Required in the POST body and optional in the PATCH body

Name	Type	Description
new_igroups	array[vdi_on_san_new_igroups]	The list of initiator groups to create. Optional in the POST or PATCH body
protection_type	protection_type	

#### storage\_service

Name	Type	Description
name	string	The storage service of the datastore. Optional in the POST or PATCH body

#### datastore

Name	Type	Description
count	integer	The number of datastores to support. Optional in the POST or PATCH body
size	integer	The size of the datastore. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body
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. Optional in the POST body
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. Required in the POST body and optional in the PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

vsi\_on\_san

A VSI application using SAN.

Name	Type	Description
datastore	<a href="#">datastore</a>	
hypervisor	string	The name of the hypervisor hosting the application. Required in the POST body
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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">vsi_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
protection_type	<a href="#">protection_type</a>	

application

Applications

Name	Type	Description
_links	<a href="#">_links</a>	
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 snapshots. Support for snapshots was added at generation 2. Any future generation numbers and their feature set will be documented
maxdata_on_san	<a href="#">maxdata_on_san</a>	MAX Data application using SAN.
mongo_db_on_san	<a href="#">mongo_db_on_san</a>	MongoDB using SAN.
name	string	Application Name. This field is user supplied when the application is created. Required in the POST body
nas	<a href="#">nas</a>	A generic NAS application.
oracle_on_nfs	<a href="#">oracle_on_nfs</a>	Oracle using NFS.
oracle_on_san	<a href="#">oracle_on_san</a>	Oracle using SAN.
oracle_rac_on_nfs	<a href="#">oracle_rac_on_nfs</a>	Oracle RAC using NFS.
oracle_rac_on_san	<a href="#">oracle_rac_on_san</a>	Oracle RAC using SAN.
protection_granularity	string	Protection granularity determines the scope of Snapshot operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot operations are performed on the entire application. If the value is "component", Snapshot operations are performed separately on the application components
rpo	<a href="#">rpo</a>	
san	<a href="#">san</a>	A generic SAN application.

Name	Type	Description
sql_on_san	<a href="#">sql_on_san</a>	Microsoft SQL using SAN.
sql_on_smb	<a href="#">sql_on_smb</a>	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	<a href="#">statistics</a>	
svm	<a href="#">svm</a>	
template	<a href="#">template</a>	
uuid	string	Application UUID. This field is generated when the application is created. Required in the URL
vdi_on_nas	<a href="#">vdi_on_nas</a>	A VDI application using NAS.
vdi_on_san	<a href="#">vdi_on_san</a>	A VDI application using SAN.
vsi_on_nas	<a href="#">vsi_on_nas</a>	A VSI application using NAS.
vsi_on_san	<a href="#">vsi_on_san</a>	A VSI application using SAN.

#### \_links

Name	Type	Description
self	<a href="#">href</a>	

#### job\_link

Name	Type	Description
<u>_links</u>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve application templates

GET /application/templates

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

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

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
num_records	integer	Number of records
records	array[ <a href="#">application_template</a> ]	



## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "description": "string",
    "maxdata_on_san": {
      "app_type": "mongodb",
      "application_components": {
        "file_system": "generic",
        "host_management_url": "string",
        "metadata": {
        },
        "protection_type": {
          "local_rpo": "6_hourly",
          "remote_rpo": "6_hourly"
        },
        "storage_service": {
          "name": "extreme"
        }
      },
      "metadata": {
      },
      "new_igroups": {
        "initiators": {
        },
        "os_type": "aix",
        "protocol": "fc"
      },
      "ocsm_url": "string",
      "os_type": "aix"
    },
    "missing_prerequisites": "string",
    "mongo_db_on_san": {
```

```

"dataset": {
  "storage_service": {
    "name": "extreme"
  }
},
"new_igroups": {
  "initiators": {
  },
  "os_type": "hyper_v",
  "protocol": "fcp"
},
"os_type": "hyper_v",
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
},
"secondary_igroups": {
}
},
"name": "string",
"nas": {
  "application_components": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "cifs_access": {
    "access": "change"
  },
  "nfs_access": {
    "access": "none"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"oracle_on_nfs": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  }
}

```

```

    }
  },
  "nfs_access": {
    "access": "none"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"oracle_on_san": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fc"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
}

```

```

    },
    "redo_log": {
      "storage_service": {
        "name": "extreme"
      }
    }
  },
  "oracle_rac_on_nfs": {
    "archive_log": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "db": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "grid_binary": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "nfs_access": {
      "access": "none"
    },
    "ora_home": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "oracle_crs": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "protection_type": {
      "local_rpo": "hourly",
      "remote_rpo": "none"
    },
    "redo_log": {
      "storage_service": {
        "name": "extreme"
      }
    }
  },
},

```

```

"oracle_rac_on_san": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db_sids": {
  },
  "grid_binary": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fcp"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_crs": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"protocol": "nas",

```

```

"san": {
  "application_components": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fc"
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"sql_on_san": {
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "hyper_v",
    "protocol": "fc"
  },
  "os_type": "windows",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "temp_db": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
},

```

```
"sql_on_smb": {
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "temp_db": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"vdi_on_nas": {
  "desktops": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"vdi_on_san": {
  "desktops": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "hypervisor": "hyper_v",
  "new_igroups": {
    "initiators": {
    },
    "protocol": "fc"
  }
}
```

```
    },
    "protection_type": {
      "local_rpo": "hourly",
      "remote_rpo": "none"
    }
  },
  "vsi_on_nas": {
    "datastore": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "nfs_access": {
      "access": "none"
    },
    "protection_type": {
      "local_rpo": "hourly",
      "remote_rpo": "none"
    }
  },
  "vsi_on_san": {
    "datastore": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "hypervisor": "hyper_v",
    "new_igroups": {
      "initiators": {
      },
      "protocol": "fc"
    },
    "protection_type": {
      "local_rpo": "hourly",
      "remote_rpo": "none"
    }
  }
}
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

self\_link

Name	Type	Description
self	<a href="#">href</a>	

metadata

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

protection\_type

Name	Type	Description
local_rpo	string	The local rpo of the application component. Optional in the POST or PATCH body
remote_rpo	string	The remote rpo of the application component. Optional in the POST or PATCH body

storage\_service

Name	Type	Description
name	string	The storage service of the application component. Optional in the POST or PATCH body

maxdata\_on\_san\_application\_components

application-components

Name	Type	Description
file_system	string	Defines the kind of file system that will be installed on this application component. Optional in the POST body
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. Required in the POST body
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. Required in the POST body and optional in the PATCH body
lun_count	integer	The number of LUNs in the application component. Required in the POST body
metadata	array[ <a href="#">metadata</a> ]	
name	string	The name of the application component. Required in the POST body and optional in the PATCH body
protection_type	<a href="#">protection_type</a>	
storage_service	<a href="#">storage_service</a>	

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

#### metadata

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

#### 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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

#### 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. Required in the POST body
application_components	array[ <a href="#">maxdata_on_san_application_components</a> ]	application-components. Optional in the POST or PATCH body
metadata	array[ <a href="#">metadata</a> ]	
new_igroups	array[ <a href="#">maxdata_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
ocsm_url	string	The OnCommand System Manager URL for this application
os_type	string	The name of the host OS running the application. Required in the POST body

#### storage\_service

Name	Type	Description
name	string	The storage service of the database. Optional in the POST or PATCH body

#### 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. Optional in the POST body
replication_factor	integer	The number of data bearing members of the replicaset, including 1 primary and at least 1 secondary. Optional in the POST body

Name	Type	Description
size	integer	The size of the database. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

protection\_type

Name	Type	Description
local_rpo	string	The local rpo of the application. Optional in the POST or PATCH body
remote_rpo	string	The remote rpo of the application. Optional in the POST body

secondary\_igroups

Name	Type	Description
name	string	The name of the initiator group for each secondary. Optional in the POST or PATCH body

## mongo\_db\_on\_san

MongoDB using SAN.

Name	Type	Description
dataset	<a href="#">dataset</a>	
new_igroups	array[ <a href="#">mongo_db_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Optional in the POST body
primary_igroup_name	string	The initiator group for the primary. Required in the POST body and optional in the PATCH body
protection_type	<a href="#">protection_type</a>	
secondary_igroups	array[ <a href="#">secondary_igroups</a> ]	

## application\_components

Name	Type	Description
name	string	The name of the application component. Optional in the POST or PATCH body
share_count	integer	The number of shares in the application component. Optional in the POST body
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member shares. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

## app\_cifs\_access

The list of CIFS access controls.

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

app\_nfs\_access

The list of NFS access controls.

Name	Type	Description
access	string	The NFS access granted. Optional in the POST body
host	string	The name of the NFS entity granted access. Optional in the POST body

nas

A generic NAS application.

Name	Type	Description
application_components	array[ <a href="#">application_components</a> ]	
cifs_access	array[ <a href="#">app_cifs_access</a> ]	The list of CIFS access controls. Optional in the POST body
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the archive log. Optional in the POST or PATCH body

archive\_log



Name	Type	Description
size	integer	The size of the archive log. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

db

Name	Type	Description
size	integer	The size of the database. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the ORACLE_HOME storage volume. Optional in the POST or PATCH body

ora\_home

Name	Type	Description
size	integer	The size of the ORACLE_HOME storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the redo log group. Optional in the POST or PATCH body

redo\_log

Name	Type	Description
mirrored	boolean	Should the redo log group be mirrored? Optional in the POST body
size	integer	The size of the redo log group. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

oracle\_on\_nfs

Oracle using NFS.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
ora_home	<a href="#">ora_home</a>	
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body

Name	Type	Description
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

oracle\_on\_san

Oracle using SAN.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">oracle_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
ora_home	<a href="#">ora_home</a>	
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the Oracle grid binary storage volume. Optional in the POST or PATCH body

grid\_binary

Name	Type	Description
size	integer	The size of the Oracle grid binary storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the Oracle CRS volume. Optional in the POST or PATCH body

oracle\_crs

Name	Type	Description
copies	integer	The number of CRS volumes. Optional in the POST body
size	integer	The size of the Oracle CRS/voting storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST body
storage_service	<a href="#">storage_service</a>	

oracle\_rac\_on\_nfs

Oracle RAC using NFS.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
grid_binary	<a href="#">grid_binary</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
ora_home	<a href="#">ora_home</a>	
oracle_crs	<a href="#">oracle_crs</a>	
protection_type	<a href="#">protection_type</a>	

Name	Type	Description
redo_log	<a href="#">redo_log</a>	

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. Optional in the POST or PATCH body

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

oracle\_rac\_on\_san

Oracle RAC using SAN.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	

Name	Type	Description
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. Optional in the POST or PATCH body
ora_home	ora_home	
oracle_crs	oracle_crs	
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	protection_type	
redo_log	redo_log	

#### 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. Optional in the POST or PATCH body
lun_count	integer	The number of LUNs in the application component. Optional in the POST body
name	string	The name of the application component. Optional in the POST or PATCH body
storage_service	storage_service	
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

## 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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

## san

A generic SAN application.

Name	Type	Description
application_components	array[ <a href="#">application_components</a> ]	
new_igroups	array[ <a href="#">san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	

## storage\_service

Name	Type	Description
name	string	The storage service of the db. Optional in the POST or PATCH body

## db

Name	Type	Description
size	integer	The size of the db. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the log db. Optional in the POST or PATCH body

log

Name	Type	Description
size	integer	The size of the log db. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body



### storage\_service

Name	Type	Description
name	string	The storage service of the temp db. Optional in the POST or PATCH body

### temp\_db

Name	Type	Description
size	integer	The size of the temp db. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

### sql\_on\_san

Microsoft SQL using SAN.

Name	Type	Description
db	<a href="#">db</a>	
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. Required in the POST body and optional in the PATCH body
log	<a href="#">log</a>	
new_igroups	array[ <a href="#">sql_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

Name	Type	Description
server_cores_count	integer	The number of server cores for the db. Optional in the POST body
temp_db	<a href="#">temp_db</a>	

#### access

Name	Type	Description
installer	string	SQL installer admin user name. Optional in the POST body
service_account	string	SQL service account user name. Required in the POST body

#### sql\_on\_smb

Microsoft SQL using SMB.

Name	Type	Description
access	<a href="#">access</a>	
db	<a href="#">db</a>	
log	<a href="#">log</a>	
protection_type	<a href="#">protection_type</a>	
server_cores_count	integer	The number of server cores for the db. Optional in the POST body
temp_db	<a href="#">temp_db</a>	

#### storage\_service

Name	Type	Description
name	string	The storage service of the desktops. Optional in the POST or PATCH body

#### desktops

Name	Type	Description
count	integer	The number of desktops to support. Optional in the POST or PATCH body
size	integer	The size of the desktops. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body
storage_service	<a href="#">storage_service</a>	

#### hyper\_v\_access

Name	Type	Description
service_account	string	Hyper-V service account. Optional in the POST body

#### vdi\_on\_nas

A VDI application using NAS.

Name	Type	Description
desktops	<a href="#">desktops</a>	
hyper_v_access	<a href="#">hyper_v_access</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

#### 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. Required in the POST body and optional in the PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

#### vdi\_on\_san

A VDI application using SAN.

Name	Type	Description
desktops	<a href="#">desktops</a>	
hypervisor	string	The name of the hypervisor hosting the application. Required in the POST body
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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">vdi_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the datastore. Optional in the POST or PATCH body

datastore

Name	Type	Description
count	integer	The number of datastores to support. Optional in the POST or PATCH body
size	integer	The size of the datastore. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body
storage_service	<a href="#">storage_service</a>	

vsi\_on\_nas

A VSI application using NAS.

Name	Type	Description
datastore	<a href="#">datastore</a>	
hyper_v_access	<a href="#">hyper_v_access</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

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. Required in the POST body and optional in the PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

vsi\_on\_san

A VSI application using SAN.

Name	Type	Description
datastore	<a href="#">datastore</a>	
hypervisor	string	The name of the hypervisor hosting the application. Required in the POST body
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. Required in the POST body and optional in the PATCH body

Name	Type	Description
new_igroups	array[ <a href="#">vsi_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
protection_type	<a href="#">protection_type</a>	

application\_template

Application Templates

Name	Type	Description
_links	<a href="#">self_link</a>	
description	string	Description
maxdata_on_san	<a href="#">maxdata_on_san</a>	MAX Data application using SAN.
missing_prerequisites	string	Missing Prerequisites
mongo_db_on_san	<a href="#">mongo_db_on_san</a>	MongoDB using SAN.
name	string	Template Name. Required in the URL
nas	<a href="#">nas</a>	A generic NAS application.
oracle_on_nfs	<a href="#">oracle_on_nfs</a>	Oracle using NFS.
oracle_on_san	<a href="#">oracle_on_san</a>	Oracle using SAN.
oracle_rac_on_nfs	<a href="#">oracle_rac_on_nfs</a>	Oracle RAC using NFS.
oracle_rac_on_san	<a href="#">oracle_rac_on_san</a>	Oracle RAC using SAN.
protocol	string	Access Protocol
san	<a href="#">san</a>	A generic SAN application.
sql_on_san	<a href="#">sql_on_san</a>	Microsoft SQL using SAN.
sql_on_smb	<a href="#">sql_on_smb</a>	Microsoft SQL using SMB.
vdi_on_nas	<a href="#">vdi_on_nas</a>	A VDI application using NAS.

Name	Type	Description
vdi_on_san	<a href="#">vdi_on_san</a>	A VDI application using SAN.
vsi_on_nas	<a href="#">vsi_on_nas</a>	A VSI application using NAS.
vsi_on_san	<a href="#">vsi_on_san</a>	A VSI application using SAN.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}

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
<code>_links</code>	<a href="#">self_link</a>	
<code>description</code>	string	Description
<code>maxdata_on_san</code>	<a href="#">maxdata_on_san</a>	MAX Data application using SAN.
<code>missing_prerequisites</code>	string	Missing Prerequisites
<code>mongo_db_on_san</code>	<a href="#">mongo_db_on_san</a>	MongoDB using SAN.
<code>name</code>	string	Template Name. Required in the URL
<code>nas</code>	<a href="#">nas</a>	A generic NAS application.
<code>oracle_on_nfs</code>	<a href="#">oracle_on_nfs</a>	Oracle using NFS.
<code>oracle_on_san</code>	<a href="#">oracle_on_san</a>	Oracle using SAN.
<code>oracle_rac_on_nfs</code>	<a href="#">oracle_rac_on_nfs</a>	Oracle RAC using NFS.
<code>oracle_rac_on_san</code>	<a href="#">oracle_rac_on_san</a>	Oracle RAC using SAN.



Name	Type	Description
protocol	string	Access Protocol
san	<a href="#">san</a>	A generic SAN application.
sql_on_san	<a href="#">sql_on_san</a>	Microsoft SQL using SAN.
sql_on_smb	<a href="#">sql_on_smb</a>	Microsoft SQL using SMB.
vdi_on_nas	<a href="#">vdi_on_nas</a>	A VDI application using NAS.
vdi_on_san	<a href="#">vdi_on_san</a>	A VDI application using SAN.
vsi_on_nas	<a href="#">vsi_on_nas</a>	A VSI application using NAS.
vsi_on_san	<a href="#">vsi_on_san</a>	A VSI application using SAN.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "description": "string",
  "maxdata_on_san": {
    "app_type": "mongodb",
    "application_components": {
      "file_system": "generic",
      "host_management_url": "string",
      "metadata": {
      },
      "protection_type": {
        "local_rpo": "6_hourly",
        "remote_rpo": "6_hourly"
      },
      "storage_service": {
        "name": "extreme"
      }
    },
    "metadata": {
    },
    "new_igroups": {
      "initiators": {
      },
      "os_type": "aix",
      "protocol": "fcp"
    },
    "ocsm_url": "string",
    "os_type": "aix"
  },
  "missing_prerequisites": "string",
  "mongo_db_on_san": {
    "dataset": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "new_igroups": {
      "initiators": {
      },
      "os_type": "hyper_v",
```

```

    "protocol": "fcp"
  },
  "os_type": "hyper_v",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "secondary_igroups": {
  }
},
"name": "string",
"nas": {
  "application_components": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "cifs_access": {
    "access": "change"
  },
  "nfs_access": {
    "access": "none"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"oracle_on_nfs": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  }
}

```

```

    },
    "protection_type": {
      "local_rpo": "hourly",
      "remote_rpo": "none"
    },
    "redo_log": {
      "storage_service": {
        "name": "extreme"
      }
    }
  },
  "oracle_on_san": {
    "archive_log": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "db": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "new_igroups": {
      "initiators": {
      },
      "os_type": "aix",
      "protocol": "fc"
    },
    "ora_home": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "os_type": "aix",
    "protection_type": {
      "local_rpo": "hourly",
      "remote_rpo": "none"
    },
    "redo_log": {
      "storage_service": {
        "name": "extreme"
      }
    }
  },
  "oracle_rac_on_nfs": {
    "archive_log": {

```

```

    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "grid_binary": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_crs": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"oracle_rac_on_san": {
  "archive_log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  }
}

```

```

    }
  },
  "db_sids": {
  },
  "grid_binary": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "aix",
    "protocol": "fcp"
  },
  "ora_home": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "oracle_crs": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "redo_log": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"protocol": "nas",
"san": {
  "application_components": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },

```

```

    "os_type": "aix",
    "protocol": "fc"
  },
  "os_type": "aix",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"sql_on_san": {
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "log": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "new_igroups": {
    "initiators": {
    },
    "os_type": "hyper_v",
    "protocol": "fc"
  },
  "os_type": "windows",
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "temp_db": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"sql_on_smb": {
  "db": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "log": {
    "storage_service": {
      "name": "extreme"
    }
  }
}

```

```

    }
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  },
  "temp_db": {
    "storage_service": {
      "name": "extreme"
    }
  }
},
"vdi_on_nas": {
  "desktops": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "nfs_access": {
    "access": "none"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"vdi_on_san": {
  "desktops": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "hypervisor": "hyper_v",
  "new_igroups": {
    "initiators": {
    },
    "protocol": "fcp"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
},
"vsi_on_nas": {
  "datastore": {
    "storage_service": {

```



```

    "name": "extreme"
  }
},
"nfs_access": {
  "access": "none"
},
"protection_type": {
  "local_rpo": "hourly",
  "remote_rpo": "none"
}
},
"vsi_on_san": {
  "datastore": {
    "storage_service": {
      "name": "extreme"
    }
  },
  "hypervisor": "hyper_v",
  "new_igroups": {
    "initiators": {
    },
    "protocol": "fcp"
  },
  "protection_type": {
    "local_rpo": "hourly",
    "remote_rpo": "none"
  }
}
}
}

```

## 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	<a href="#">href</a>	

metadata

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

protection\_type

Name	Type	Description
local_rpo	string	The local rpo of the application component. Optional in the POST or PATCH body
remote_rpo	string	The remote rpo of the application component. Optional in the POST or PATCH body

storage\_service

Name	Type	Description
name	string	The storage service of the application component. Optional in the POST or PATCH body

maxdata\_on\_san\_application\_components

application-components

Name	Type	Description
file_system	string	Defines the kind of file system that will be installed on this application component. Optional in the POST body
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. Required in the POST body
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. Required in the POST body and optional in the PATCH body
lun_count	integer	The number of LUNs in the application component. Required in the POST body
metadata	array[ <a href="#">metadata</a> ]	
name	string	The name of the application component. Required in the POST body and optional in the PATCH body
protection_type	<a href="#">protection_type</a>	
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body

metadata

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

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

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. Required in the POST body
application_components	array[ <a href="#">maxdata_on_san_application_components</a> ]	application-components. Optional in the POST or PATCH body
metadata	array[ <a href="#">metadata</a> ]	
new_igroups	array[ <a href="#">maxdata_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body

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. Required in the POST body

#### storage\_service

Name	Type	Description
name	string	The storage service of the database. Optional in the POST or PATCH body

#### 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. Optional in the POST body
replication_factor	integer	The number of data bearing members of the replicaset, including 1 primary and at least 1 secondary. Optional in the POST body
size	integer	The size of the database. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

#### mongo\_db\_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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

#### protection\_type

Name	Type	Description
local_rpo	string	The local rpo of the application. Optional in the POST or PATCH body
remote_rpo	string	The remote rpo of the application. Optional in the POST body

#### secondary\_igroups

Name	Type	Description
name	string	The name of the initiator group for each secondary. Optional in the POST or PATCH body

#### mongo\_db\_on\_san

MongoDB using SAN.

Name	Type	Description
dataset	<a href="#">dataset</a>	
new_igroups	array[ <a href="#">mongo_db_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body

Name	Type	Description
os_type	string	The name of the host OS running the application. Optional in the POST body
primary_igroup_name	string	The initiator group for the primary. Required in the POST body and optional in the PATCH body
protection_type	<a href="#">protection_type</a>	
secondary_igroups	array[ <a href="#">secondary_igroups</a> ]	

#### application\_components

Name	Type	Description
name	string	The name of the application component. Optional in the POST or PATCH body
share_count	integer	The number of shares in the application component. Optional in the POST body
storage_service	<a href="#">storage_service</a>	
total_size	integer	The total size of the application component, split across the member shares. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

#### app\_cifs\_access

The list of CIFS access controls.

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

#### app\_nfs\_access

The list of NFS access controls.



Name	Type	Description
access	string	The NFS access granted. Optional in the POST body
host	string	The name of the NFS entity granted access. Optional in the POST body

nas

A generic NAS application.

Name	Type	Description
application_components	array[ <a href="#">application_components</a> ]	
cifs_access	array[ <a href="#">app_cifs_access</a> ]	The list of CIFS access controls. Optional in the POST body
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the archive log. Optional in the POST or PATCH body

archive\_log

Name	Type	Description
size	integer	The size of the archive log. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

db

Name	Type	Description
size	integer	The size of the database. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the ORACLE_HOME storage volume. Optional in the POST or PATCH body

ora\_home

Name	Type	Description
size	integer	The size of the ORACLE_HOME storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the redo log group. Optional in the POST or PATCH body

redo\_log

Name	Type	Description
mirrored	boolean	Should the redo log group be mirrored? Optional in the POST body

Name	Type	Description
size	integer	The size of the redo log group. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

oracle\_on\_nfs

Oracle using NFS.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
ora_home	<a href="#">ora_home</a>	
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

oracle\_on\_san

Oracle using SAN.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">oracle_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
ora_home	<a href="#">ora_home</a>	
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the Oracle grid binary storage volume. Optional in the POST or PATCH body

grid\_binary

Name	Type	Description
size	integer	The size of the Oracle grid binary storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

Name	Type	Description
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the Oracle CRS volume. Optional in the POST or PATCH body

oracle\_crs

Name	Type	Description
copies	integer	The number of CRS volumes. Optional in the POST body
size	integer	The size of the Oracle CRS/voting storage volume. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST body
storage_service	<a href="#">storage_service</a>	

oracle\_rac\_on\_nfs

Oracle RAC using NFS.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
grid_binary	<a href="#">grid_binary</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
ora_home	<a href="#">ora_home</a>	
oracle_crs	<a href="#">oracle_crs</a>	
protection_type	<a href="#">protection_type</a>	
redo_log	<a href="#">redo_log</a>	

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. Optional in the POST or PATCH body

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

oracle\_rac\_on\_san

Oracle RAC using SAN.

Name	Type	Description
archive_log	<a href="#">archive_log</a>	
db	<a href="#">db</a>	
db_sids	array[ <a href="#">db_sids</a> ]	
grid_binary	<a href="#">grid_binary</a>	

Name	Type	Description
new_igroups	array[oracle_rac_on_san_new_igroups]	The list of initiator groups to create. Optional in the POST or PATCH body
ora_home	ora_home	
oracle_crs	oracle_crs	
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	protection_type	
redo_log	redo_log	

#### 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. Optional in the POST or PATCH body
lun_count	integer	The number of LUNs in the application component. Optional in the POST body
name	string	The name of the application component. Optional in the POST or PATCH body
storage_service	storage_service	
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body

#### 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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

san

A generic SAN application.

Name	Type	Description
application_components	array[ <a href="#">application_components</a> ]	
new_igroups	array[ <a href="#">san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Required in the POST body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the db. Optional in the POST or PATCH body

db



Name	Type	Description
size	integer	The size of the db. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the log db. Optional in the POST or PATCH body

log

Name	Type	Description
size	integer	The size of the log db. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body and optional in the PATCH body
storage_service	<a href="#">storage_service</a>	

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. Required in the POST body and optional in the PATCH body
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application. Optional in the POST or PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

### storage\_service

Name	Type	Description
name	string	The storage service of the temp db. Optional in the POST or PATCH body

### temp\_db

Name	Type	Description
size	integer	The size of the temp db. Usage: {<integer>[KB MB GB TB PB]} Optional in the POST or PATCH body
storage_service	<a href="#">storage_service</a>	

### sql\_on\_san

Microsoft SQL using SAN.

Name	Type	Description
db	<a href="#">db</a>	
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. Required in the POST body and optional in the PATCH body
log	<a href="#">log</a>	
new_igroups	array[ <a href="#">sql_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
os_type	string	The name of the host OS running the application. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

Name	Type	Description
server_cores_count	integer	The number of server cores for the db. Optional in the POST body
temp_db	<a href="#">temp_db</a>	

#### access

Name	Type	Description
installer	string	SQL installer admin user name. Optional in the POST body
service_account	string	SQL service account user name. Required in the POST body

#### sql\_on\_smb

Microsoft SQL using SMB.

Name	Type	Description
access	<a href="#">access</a>	
db	<a href="#">db</a>	
log	<a href="#">log</a>	
protection_type	<a href="#">protection_type</a>	
server_cores_count	integer	The number of server cores for the db. Optional in the POST body
temp_db	<a href="#">temp_db</a>	

#### storage\_service

Name	Type	Description
name	string	The storage service of the desktops. Optional in the POST or PATCH body

#### desktops

Name	Type	Description
count	integer	The number of desktops to support. Optional in the POST or PATCH body
size	integer	The size of the desktops. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body
storage_service	<a href="#">storage_service</a>	

#### hyper\_v\_access

Name	Type	Description
service_account	string	Hyper-V service account. Optional in the POST body

#### vdi\_on\_nas

A VDI application using NAS.

Name	Type	Description
desktops	<a href="#">desktops</a>	
hyper_v_access	<a href="#">hyper_v_access</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

#### 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. Required in the POST body and optional in the PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

#### vdi\_on\_san

A VDI application using SAN.

Name	Type	Description
desktops	<a href="#">desktops</a>	
hypervisor	string	The name of the hypervisor hosting the application. Required in the POST body
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. Required in the POST body and optional in the PATCH body
new_igroups	array[ <a href="#">vdi_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
protection_type	<a href="#">protection_type</a>	

storage\_service

Name	Type	Description
name	string	The storage service of the datastore. Optional in the POST or PATCH body

datastore

Name	Type	Description
count	integer	The number of datastores to support. Optional in the POST or PATCH body
size	integer	The size of the datastore. Usage: {<integer>[KB MB GB TB PB]} Required in the POST body
storage_service	<a href="#">storage_service</a>	

vsi\_on\_nas

A VSI application using NAS.

Name	Type	Description
datastore	<a href="#">datastore</a>	
hyper_v_access	<a href="#">hyper_v_access</a>	
nfs_access	array[ <a href="#">app_nfs_access</a> ]	The list of NFS access controls. Optional in the POST body
protection_type	<a href="#">protection_type</a>	

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. Required in the POST body and optional in the PATCH body
protocol	string	The protocol of the new initiator group. Optional in the POST or PATCH body

vsi\_on\_san

A VSI application using SAN.

Name	Type	Description
datastore	<a href="#">datastore</a>	
hypervisor	string	The name of the hypervisor hosting the application. Required in the POST body
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. Required in the POST body and optional in the PATCH body

Name	Type	Description
new_igroups	array[ <a href="#">vsi_on_san_new_igroups</a> ]	The list of initiator groups to create. Optional in the POST or PATCH body
protection_type	<a href="#">protection_type</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

Retrieves the collection of cloud targets in the cluster.

### Related ONTAP commands

- `storage aggregate object-store config show`

### Learn more

- [DOC /cloud/targets](#)

### Parameters

Name	Type	In	Required	Description
ssl_enabled	boolean	query	False	Filter by ssl_enabled
port	integer	query	False	Filter by port
used	integer	query	False	Filter by used
owner	string	query	False	Filter by owner
authentication_type	string	query	False	Filter by authentication_type
certificate_validation_enabled	boolean	query	False	Filter by certificate_validation_enabled
server	string	query	False	Filter by server



<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
cap_url	string	query	False	Filter by cap_url
snapmirror_use	string	query	False	Filter by snapmirror_use
uuid	string	query	False	Filter by uuid
name	string	query	False	Filter by name
provider_type	string	query	False	Filter by provider_type
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
azure_account	string	query	False	Filter by azure_account
ipspace.uuid	string	query	False	Filter by ipspace.uuid
ipspace.name	string	query	False	Filter by ipspace.name
container	string	query	False	Filter by container
access_key	string	query	False	Filter by access_key
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.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">cloud_target</a> ]	

**Example response**

A large, empty rectangular box with a thin, dashed border, occupying most of the page. It is intended for an example response.

```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "authentication_type": "key",
    "cap_url":
    "https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency&mission=mymission&role=myrole",
    "container": "bucket1",
    "ipspace": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "exchange",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "owner": "fabricpool",
    "snapmirror_use": "data",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "used": 0,
    "uuid": "string"
  }
}

```

## Error

Status: Default, n/a

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

ipspace

IPspace to use in order to reach the cloud target.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cloud\_target

Name	Type	Description
_links	<a href="#">_links</a>	

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 and IBM_COS provider types.
container	string	Data bucket/container name
ipspace	<a href="#">ipspace</a>	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.

Name	Type	Description
provider_type	string	Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoggleCloud, IBM_COS, and SGWS are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS.
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>
snapmirror_use	string	Use of the cloud target by SnapMirror.
ssl_enabled	boolean	SSL/HTTPS enabled or not
svm	<a href="#">svm</a>	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.
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.



Name	Type	Description
uuid	string	Cloud target UUID

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

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*.
- `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*.
- `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.

### Default property values

- `authentication_type`
- *ec2\_iam* - if running in Cloud Volumes ONTAP in AWS
- *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*

### Related ONTAP commands

- `storage aggregate object-store config create`

### Learn more

- [DOC /cloud/targets](#)

## Parameters

Name	Type	In	Required	Description
ignore_warnings	boolean	query	False	Specifies whether or not warning codes should be ignored.
check_only	boolean	query	False	Do not create the target configuration, only check that the POST request succeeds.

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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-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 and IBM_COS provider types.
container	string	Data bucket/container name
ipspace	<a href="#">ipspace</a>	IPspace to use in order to reach the cloud target.

Name	Type	Description
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.
provider_type	string	Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoggleCloud, IBM_COS, and SGWS are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS.
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>
snapmirror_use	string	Use of the cloud target by SnapMirror.
ssl_enabled	boolean	SSL/HTTPS enabled or not

Name	Type	Description
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.
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

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_type": "key",
  "cap_url":
  "https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency&mission=mymission&role=myrole",
  "container": "bucket1",
  "ipSPACE": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "exchange",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"owner": "fabricpool",
"snapmirror_use": "data",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"name": "svm1",
"uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"used": 0,
"uuid": "string"
}
```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

ipspace

IPspace to use in order to reach the cloud target.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cloud\_target

Name	Type	Description
_links	<a href="#">_links</a>	
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-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 and IBM_COS provider types.
container	string	Data bucket/container name
ipspace	<a href="#">ipspace</a>	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.
provider_type	string	Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoggleCloud, IBM_COS, and SGWS are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS.

Name	Type	Description
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>
snapmirror_use	string	Use of the cloud target by SnapMirror.
ssl_enabled	boolean	SSL/HTTPS enabled or not
svm	<a href="#">svm</a>	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.
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
<a href="#">_links</a>	<a href="#">_links</a>	

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[ <a href="#">error_arguments</a> ]	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}

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`

### Learn more

- [DOC /cloud/targets](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

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

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

Retrieves the cloud target specified by the UUID.

### Related ONTAP commands

- `storage aggregate object-store config show`

### Learn more

- [DOC /cloud/targets](#)

### 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
<a href="#">_links</a>	<a href="#">_links</a>	
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-password, and session token) for accessing the object store.

Name	Type	Description
certificate_validation_enabled	boolean	Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS and IBM_COS provider types.
container	string	Data bucket/container name
ipspace	<a href="#">ipspace</a>	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.
provider_type	string	Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoggleCloud, IBM_COS, and SGWS are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS.
secret_password	string	Secret access key for AWS_S3 and other S3 compatible provider types.

Name	Type	Description
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>
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.
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"
    }
  },
  "authentication_type": "key",
  "cap_url":
  "https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency&mission=mymission&role=myrole",
  "container": "bucket1",
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "exchange",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"owner": "fabricpool",
"snapmirror_use": "data",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"name": "svm1",
"uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"used": 0,
"uuid": "string"
}
```

## Error

Status: Default, n/a

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	<a href="#">href</a>	

ipspace

IPspace to use in order to reach the cloud target.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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`

### Learn more

- [DOC /cloud/targets](#)

### 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.
check_only	boolean	query	False	Do not modify the configuration, only check that the PATCH request succeeds.

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	

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-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 and IBM_COS provider types.
container	string	Data bucket/container name
ipspace	<a href="#">ipspace</a>	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.

Name	Type	Description
provider_type	string	Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoggleCloud, IBM_COS, and SGWS are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS.
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>
snapmirror_use	string	Use of the cloud target by SnapMirror.
ssl_enabled	boolean	SSL/HTTPS enabled or not
svm	<a href="#">svm</a>	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.
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

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_type": "key",
  "cap_url":
  "https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency&mission=mymission&role=myrole",
  "container": "bucket1",
  "ipSPACE": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "exchange",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"owner": "fabricpool",
"snapmirror_use": "data",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"name": "svm1",
"uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"used": 0,
"uuid": "string"
}
```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

ipspace

IPspace to use in order to reach the cloud target.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cloud\_target

Name	Type	Description
_links	<a href="#">_links</a>	
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-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 and IBM_COS provider types.
container	string	Data bucket/container name
ipspace	<a href="#">ipspace</a>	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.
provider_type	string	Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoggleCloud, IBM_COS, and SGWS are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS.

Name	Type	Description
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>
snapmirror_use	string	Use of the cloud target by SnapMirror.
ssl_enabled	boolean	SSL/HTTPS enabled or not
svm	<a href="#">svm</a>	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.
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
<a href="#">_links</a>	<a href="#">_links</a>	

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[ <a href="#">error_arguments</a> ]	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.

## APIs

---

### cluster

The cluster APIs cover basic management of the cluster, including viewing and modifying settings such as the name, UUID, version, NTP servers, DNS domains, and the nodes in the cluster. POST /api/cluster allows creation of the cluster, including adding all of the nodes to the cluster available during setup, initial configuration of cluster and node management interfaces, NTP servers, name servers, licenses, and node names.

---

### nodes

The node APIs allow you to gather information about the nodes in a cluster, including model number, serial number, HA Group information, names, management interfaces, and UUIDs. By setting a query parameter, the administrator can also discover nodes on the cluster network that have not been added to the cluster and use that information in the POST operation to add them. Nodes can be added to the cluster through the POST API.

---

### jobs

The job APIs are used to monitor the progress of running and recent jobs using GET. Some jobs are capable of being paused and cancelled using the PATCH operation.

---

## Manage clusters

### Cluster endpoint overview

#### Overview

This API is used 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

The following fields are always required for any POST /cluster request:

- name
- password

## Optional cluster-wide configuration

The following fields are used to setup additional cluster-wide configuration:

- location
- contact
- dns\_domains
- name\_servers
- ntp\_servers
- license
- configuration\_backup
- management\_interface
- nodes

## Nodes field

The nodes field specifies the nodes to join to the cluster. All nodes must be at the same version to use this API. If no nodes are specified, the cluster is configured with one node added. The node added is the node to which the REST request is issued. If one node is specified, the "node.cluster\_interface.ip.address" field must not be used. If multiple nodes are specified, the node to which the REST request is issued must be provided in addition to the remote nodes, and the "node.cluster\_interface.ip.address" field is required for each node to identify them. All other node fields are optional in all cases. If a field is provided for one node, it must be provided for all nodes.

## Node networking fields

The cluster management interface and each node management interface use the cluster management interface netmask and gateway. For advanced configurations where the cluster and node management interfaces are on different subnets, the /network/ip/interface APIs must be used 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. A node field for this node can be provided for node-specific configuration and the "node.cluster\_interface.ip.address" field must not be used. 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.

## Performance monitoring

Performance of the cluster can be monitored by the `metric.*` and `statistics.*` fields. These show the performance of the cluster 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.

## 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` has been 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 can be monitored with the job UUID using the `/cluster/jobs` API. The "message" field in the response of GET `/cluster/jobs/{uuid}` 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 in the 'Responses' of the POST `/cluster` documentation for common error codes and descriptions.

### Re-running POST `/cluster`

The POST `/cluster` command can be re-run if errors occur. When re-running the request, the same body and query parameters must be used. The value of any field in the original body or query can be changed, but the fields that were provided cannot be changed. 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 re-run 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 re-run request with the following body details is invalid:

+

```
body =
{
  "name": "clusCreateRerun",
  "password": "openSesame",
  "nodes": [
    {
      "cluster_interface": {
        "ip": {
          "address": "3.3.3.3"
        }
      }
    }
  ]
}
```

Also, note that the password might already be configured. If a password is already configured and a new password is provided, this request overwrites the existing password. If a password is already configured either by another interface or by a previous POST to /cluster, any future REST requests must be authenticated with that password. If POST to /cluster with the default return\_timeout of 0 returns an error, then the password was not changed.

### **Connection failures during cluster create**

There are two cases where a request to poll the job status might fail during the cluster create job. In these cases, programmatic use of the RESTful interface should 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, return errors on GET, or connection timeouts might occur. Any 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

### **Examples**

#### **A minimal configuration of a 2-node setup**

---

```
# Body
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 body
```

---

## A single node setup with additional node configuration

---

```
# Body
body =
{
"name": "clusCreateExample2",
"password": "openSesame",
"nodes": [
  {
    "name": "singleNode",
    "location": "Sunnyvale"
  }
]
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster?single_node_cluster=true" -d
body
```

---

## Modifying a cluster-wide configuration

```
# Body
body =
{
"contact": "it@company.com"
}

# Request
curl -X PATCH "https://<mgmt-ip>/api/cluster" -d body
```

---

## A detailed example of a cluster "create" operation

The following is an example of how a cluster can be created using the cluster APIs. 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.

### 1) 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 network, as well as the connections to the management network, must be completed. Once the nodes are powered up, 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 up 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.
```

## 2) Discovering the nodes

Issuing 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 network. Information returned include 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 those nodes that should not be part of the cluster. If not enough nodes are discovered, ensure all the nodes are powered up, that the connections to the cluster network are complete, and retry the command.

+

```
# The API:
/api/cluster/nodes

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/nodes?fields=*" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "60277d87-19e4-11e9-ba25-005056bb6eee",
      "name": "Computer.local",
      "serial_number": "4136233-26-3",
      "model": "FAS9000",
      "version": {
        "full": "NetApp Release 9.6.0: Wed Jan 16 18:20:57 UTC 2019",
        "generation": 9,
        "major": 6,
        "minor": 0
      },
      "membership": "available",
      "cluster_interfaces": [
        {
          "ip": {
            "address": "169.254.245.113"
          }
        }
      ],
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/60277d87-19e4-11e9-ba25-005056bb6eee"
        }
      }
    },
    {
      "uuid": "8071ba1b-19e3-11e9-b003-005056bb096a",
      "name": "Computer-6.local",
      "serial_number": "4136233-26-2",
      "model": "FAS9000",
```

```

"version": {
  "full": "NetApp Release 9.6.0: Wed Jan 16 18:20:57 UTC 2019",
  "generation": 9,
  "major": 6,
  "minor": 0
},
"membership": "available",
"cluster_interfaces": [
  {
    "ip": {
      "address": "169.254.217.95"
    }
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/nodes/8071ba1b-19e3-11e9-b003-005056bb096a"
  }
}
},
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/cluster/nodes?fields=*"
  }
}
}

```

### 3) Creating the cluster

Once the node information is available, including each node's cluster interface address, the information for creating the cluster can be assembled. You must 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. Each node to be included in the cluster must have the cluster interface address provided so that it can be connected to while adding it to the cluster. In addition to the cluster interface address, the optional node name, location, and management interface information can be supplied. If node names are not provided, nodes are named based on the cluster name. The nodes' management interface netmask and gateway values are omitted and must be the same as the cluster management interface's netmask 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":["AMEPOS0IKLKGEEEEEDGNDEKSJDE"]}
,"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"
      }
    }
  }
}

```

#### 4) Monitoring the progress of cluster creation

To monitor the progress of the cluster create operation, the job link returned should be polled 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"
    }
  }
}
```

## 5) Verifying the cluster information

Once the cluster is created, the information applied can be verified using a number of APIs. Most of the information provided can be retrieved using the `/api/cluster` and `/api/cluster/nodes` APIs. In addition, the network interface and route information can be viewed using the `/api/network` APIs. The following example details how to retrieve the cluster information:

+

```
# The API:
/api/cluster

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster" -H "accept: application/hal+json"

# The response:
{
  "name": "cluster1",
  "uuid": "93d05f83-7d80-482d-b59c-a6661d272a47",
  "location": "datacenter1",
  "contact": "me",
}
```

```
"version": {
  "full": "NetApp Release 9.6.0: Wed Jan 16 18:20:57 UTC 2019",
  "generation": 9,
  "major": 6,
  "minor": 0
},
"dns_domains": [
  "example.com"
],
"name_servers": [
  "10.224.223.130",
  "10.224.223.131",
  "10.224.223.132"
],
"ntp_servers": [
  "time.nist.gov"
],
"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"
      }
    }
  }
],
"metric": {
  "timestamp": "2019-04-09T06:33:30Z",
  "duration": "PT15S",
  "status": "ok",
  "latency": {
    "other": 0,
    "total": 525,
    "read": 525,
    "write": 0
  },
  "iops": {
    "read": 200,
    "write": 0,
    "other": 0,
```

```
    "total": 200
  },
  "throughput": {
    "read": 820838,
    "write": 0,
    "other": 0,
    "total": 820838
  }
},
"statistics": {
  "timestamp": "2019-04-09T06:33:50Z",
  "status": "ok",
  "latency_raw": {
    "other": 38928,
    "total": 3331918704,
    "read": 3331879776,
    "write": 0
  },
  "iops_raw": {
    "read": 6188132,
    "write": 0,
    "other": 5,
    "total": 6188137
  },
  "throughput_raw": {
    "read": 25346587876,
    "write": 0,
    "other": 0,
    "total": 25346587876
  }
},
"_links": {
  "self": {
    "href": "/api/cluster"
  }
}
}
```

---

## Retrieve the cluster configuration

GET /cluster

Retrieves the cluster configuration.

## Learn more

- [DOC /cluster](#)

## Parameters

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

## Response

```
Status: 200, Ok
```

Name	Type	Description
_links	<a href="#">_links</a>	
configuration_backup	<a href="#">configuration_backup</a>	
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"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>
license	license	License keys or NLF contents.
location	string	
management_interface	management_interface	The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.
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.

Name	Type	Description
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.
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.
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"
    }
  },
  "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": {
    }
  },
  "location": "building 1",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    }
  },
  "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": "2017-01-25 11:20:13 UTC"
},
"name": "cluster1",
"name_servers": [
  "10.224.65.20",
  "2001:db08:a0b:12f0::1"
],
"nodes": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "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": 102400000000,
      "firmware_version": "NA05",
      "hardware_revision": "A1",
      "model": "X1970A",
      "part_number": "119-00207",
      "serial_number": "A22P5061550000187",
      "slot": "6-1",
      "state": "ok"
    },
    "frus": {
      "id": 0,
      "state": "ok",
      "type": "fan"
    },
    "over_temperature": "over"
  },
  "date": "2017-01-25 11:20:13 +0400",
  "ha": {
    "partners": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "location": "rack 2 row 5",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "management_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },

```

```
"ip": {
  "address": "10.10.10.7"
},
"name": "lif1",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"membership": "available",
"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": "up",
  "mac_address": "string",
  "state": "online"
},
"uptime": 300536,
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
"version": {
  "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
  "generation": 9,
  "major": 4,
  "minor": 0
}
},
"ntp_servers": [
  "time.nist.gov",
  "10.98.19.20",
  "2610:20:6F15:15::27"
],
"password": "mypassword",
"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": "2017-01-25 11:20:13 UTC"
},
"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	<a href="#">href</a>	

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	

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.

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.

#### management\_interface

The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.

Name	Type	Description
ip	<a href="#">ip</a>	Object to setup an interface along with its default router.

#### ip

##### IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

#### management\_interfaces

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

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
ip	<a href="#">ip</a>	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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">iops</a>	The rate of I/O operations observed at the storage object.
latency	<a href="#">latency</a>	The round trip latency in microseconds observed at the storage object.

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 does not have the latest data.
throughput	<a href="#">throughput</a>	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	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

## cluster\_interfaces

### Network interface

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
ip	<a href="#">ip</a>	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[ <a href="#">flash_cache</a> ]	A list of Flash-Cache devices. Only returned when requested by name.

Name	Type	Description
frus	array[frus]	A list of frus in 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.

#### partners

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	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.
partners	array[ <a href="#">partners</a> ]	The nodes in this node's High Availability (HA) group.

#### management\_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

#### management\_interfaces

Network interface

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
<a href="#">ip</a>	<a href="#">ip</a>	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

#### 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	<a href="#">ipv4_interface</a>	Object to setup an interface along with its default router.
ipv6_interface	<a href="#">ipv6_interface</a>	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

#### 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.

#### nodes

Complete node information

Name	Type	Description
_links	<a href="#">_links</a>	
cluster_interface	<a href="#">cluster_interface</a>	The cluster network IP address of the node to be added.
cluster_interfaces	array[ <a href="#">cluster_interfaces</a> ]	
controller	<a href="#">controller</a>	Controller information

Name	Type	Description
date	string	Specifies the ISO-8601 format date and time on the node.
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> <li>• <i>available</i> - If a node is available, this means it 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. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.</li> <li>• <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may 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.</li> <li>• <i>member</i> - Nodes that are members have successfully joined the cluster.</li> </ul>
model	string	
name	string	

Name	Type	Description
serial_number	string	
service_processor	<a href="#">service_processor</a>	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	<a href="#">version</a>	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.

#### 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.
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 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 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 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.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the cluster configuration

PATCH `/cluster`

Updates the cluster configuration once the cluster has been created.

Learn more

- [DOC /cluster](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>configuration_backup</code>	<a href="#">configuration_backup</a>	
<code>contact</code>	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"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>
license	<a href="#">license</a>	License keys or NLF contents.
location	string	
management_interface	<a href="#">management_interface</a>	The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.
management_interfaces	array[ <a href="#">management_interfaces</a> ]	
metric	<a href="#">metric</a>	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.

Name	Type	Description
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.
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.
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 request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "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": {
    }
  },
  "location": "building 1",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    }
  },
  "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": "2017-01-25 11:20:13 UTC"
},
"name": "cluster1",
"name_servers": [
  "10.224.65.20",
  "2001:db08:a0b:12f0::1"
],
"nodes": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "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": 102400000000,
      "firmware_version": "NA05",
      "hardware_revision": "A1",
      "model": "X1970A",
      "part_number": "119-00207",
      "serial_number": "A22P5061550000187",
      "slot": "6-1",
      "state": "ok"
    },
    "frus": {
      "id": 0,
      "state": "ok",
      "type": "fan"
    },
    "over_temperature": "over"
  },
  "date": "2017-01-25 11:20:13 +0400",
  "ha": {
    "partners": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "location": "rack 2 row 5",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "management_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },

```



```
"ip": {
  "address": "10.10.10.7"
},
"name": "lif1",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"membership": "available",
"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": "up",
  "mac_address": "string",
  "state": "online"
},
"uptime": 300536,
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
"version": {
  "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
  "generation": 9,
  "major": 4,
  "minor": 0
}
},
"ntp_servers": [
  "time.nist.gov",
  "10.98.19.20",
  "2610:20:6F15:15::27"
],
"password": "mypassword",
"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": "2017-01-25 11:20:13 UTC"
  },
  "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
  }
}

```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

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

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
8847362	Too many name servers provided.
8847361	Too many DNS domains provided.
9240587	A name must be provided.
131727388	Hostnames for NTP servers cannot be used without DNS configured.
2097165	An NTP server could not be reached.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

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	

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.

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.

#### management\_interface

The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.

Name	Type	Description
ip	<a href="#">ip</a>	Object to setup an interface along with its default router.

#### ip

##### IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

#### management\_interfaces

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

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
<a href="#">ip</a>	<a href="#">ip</a>	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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">iops</a>	The rate of I/O operations observed at the storage object.
latency	<a href="#">latency</a>	The round trip latency in microseconds observed at the storage object.



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 does not have the latest data.
throughput	<a href="#">throughput</a>	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	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

## cluster\_interfaces

### Network interface

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
ip	<a href="#">ip</a>	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[ <a href="#">flash_cache</a> ]	A list of Flash-Cache devices. Only returned when requested by name.

Name	Type	Description
frus	array[frus]	A list of frus in 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.

#### partners

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	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.
partners	array[ <a href="#">partners</a> ]	The nodes in this node's High Availability (HA) group.

#### management\_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

#### management\_interfaces

#### Network interface

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

#### 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	<a href="#">ipv4_interface</a>	Object to setup an interface along with its default router.
ipv6_interface	<a href="#">ipv6_interface</a>	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

#### 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.

#### nodes

Complete node information

Name	Type	Description
_links	<a href="#">_links</a>	
cluster_interface	<a href="#">cluster_interface</a>	The cluster network IP address of the node to be added.
cluster_interfaces	array[ <a href="#">cluster_interfaces</a> ]	
controller	<a href="#">controller</a>	Controller information

Name	Type	Description
date	string	Specifies the ISO-8601 format date and time on the node.
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> <li>• <i>available</i> - If a node is available, this means it 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. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.</li> <li>• <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may 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.</li> <li>• <i>member</i> - Nodes that are members have successfully joined the cluster.</li> </ul>
model	string	
name	string	

Name	Type	Description
serial_number	string	
service_processor	<a href="#">service_processor</a>	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	<a href="#">version</a>	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.

#### 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.
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 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	<a href="#">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	<a href="#">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	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 does not have the latest data.
throughput_raw	<a href="#">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	The timestamp of the performance data.

cluster

Complete cluster information

Name	Type	Description
_links	<a href="#">_links</a>	
configuration_backup	<a href="#">configuration_backup</a>	
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"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>
license	<a href="#">license</a>	License keys or NLF contents.
location	string	
management_interface	<a href="#">management_interface</a>	The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.
management_interfaces	array[ <a href="#">management_interfaces</a> ]	

Name	Type	Description
metric	<a href="#">metric</a>	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.
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.
statistics	<a href="#">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.
uuid	string	
version	<a href="#">version</a>	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.

#### job\_link

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create a cluster

POST /cluster

Sets up a cluster.

### Required properties

- name
- password

### Recommended optional properties

- location
- contact
- dns\_domains
- name\_servers
- ntp\_servers
- license
- configuration\_backup
- management\_interface
- nodes

## 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.

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
configuration_backup	<a href="#">configuration_backup</a>	
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"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>
license	<a href="#">license</a>	License keys or NLF contents.
location	string	
management_interface	<a href="#">management_interface</a>	The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.
management_interfaces	array[ <a href="#">management_interfaces</a> ]	
metric	<a href="#">metric</a>	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.

Name	Type	Description
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.
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.
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 request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "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": {
    }
  },
  "location": "building 1",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    }
  },
  "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": "2017-01-25 11:20:13 UTC"
},
"name": "cluster1",
"name_servers": [
  "10.224.65.20",
  "2001:db08:a0b:12f0::1"
],
"nodes": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "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": 102400000000,
      "firmware_version": "NA05",
      "hardware_revision": "A1",
      "model": "X1970A",
      "part_number": "119-00207",
      "serial_number": "A22P5061550000187",
      "slot": "6-1",
      "state": "ok"
    },
    "frus": {
      "id": 0,
      "state": "ok",
      "type": "fan"
    },
    "over_temperature": "over"
  },
  "date": "2017-01-25 11:20:13 +0400",
  "ha": {
    "partners": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "location": "rack 2 row 5",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "management_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },

```

```
"ip": {
  "address": "10.10.10.7"
},
"name": "lif1",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"membership": "available",
"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": "up",
  "mac_address": "string",
  "state": "online"
},
"uptime": 300536,
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
"version": {
  "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
  "generation": 9,
  "major": 4,
  "minor": 0
}
},
"ntp_servers": [
  "time.nist.gov",
  "10.98.19.20",
  "2610:20:6F15:15::27"
],
"password": "mypassword",
"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": "2017-01-25 11:20:13 UTC"
  },
  "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
  }
}

```

### Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

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

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
131727360	A node could not be added to the cluster. This is a generic code, see response message for details.
8978433	An invalid license key was provided.
8847362	Too many name servers provided.
8847361	Too many DNS domains provided.
9240587	A name must be provided.
39387137	The URL provided is invalid.
131727389	URL and username are required for configuration backup.
262245	The value provided is invalid.
1179817	The IP address, netmask, and gateway must all be provided for cluster management interface.
1179813	Fields set for one node must be set for all nodes.
1179818	The IP address and gateway must be of the same family.
1179821	An IP address and netmask conflicts with an existing entry.
131727388	Hostnames for NTP servers cannot be used without DNS configured.

Error Code	Description
2097165	An NTP server could not be reached.
1179825	All management and cluster config IP addresses must belong to the same address family.
8847394	An invalid DNS domain was provided.
9240594	An invalid name was provided.
1179824	An invalid gateway was provided.

Name	Type	Description
error	<a href="#">error</a>	

### 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	<a href="#">href</a>	

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	

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.

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.

#### management\_interface

The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.

Name	Type	Description
ip	<a href="#">ip</a>	Object to setup an interface along with its default router.

#### ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

#### management\_interfaces

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

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
ip	<a href="#">ip</a>	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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">iops</a>	The rate of I/O operations observed at the storage object.
latency	<a href="#">latency</a>	The round trip latency in microseconds observed at the storage object.

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 does not have the latest data.
throughput	<a href="#">throughput</a>	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	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

## cluster\_interfaces

### Network interface

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
ip	<a href="#">ip</a>	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[ <a href="#">flash_cache</a> ]	A list of Flash-Cache devices. Only returned when requested by name.

Name	Type	Description
frus	array[frus]	A list of frus in 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.

#### partners

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	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.
partners	array[ <a href="#">partners</a> ]	The nodes in this node's High Availability (HA) group.

#### management\_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

#### management\_interfaces

#### Network interface

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
<a href="#">ip</a>	<a href="#">ip</a>	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

#### 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	<a href="#">ipv4_interface</a>	Object to setup an interface along with its default router.
ipv6_interface	<a href="#">ipv6_interface</a>	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

#### 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.

#### nodes

Complete node information

Name	Type	Description
_links	<a href="#">_links</a>	
cluster_interface	<a href="#">cluster_interface</a>	The cluster network IP address of the node to be added.
cluster_interfaces	array[ <a href="#">cluster_interfaces</a> ]	
controller	<a href="#">controller</a>	Controller information

Name	Type	Description
date	string	Specifies the ISO-8601 format date and time on the node.
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> <li>• <i>available</i> - If a node is available, this means it 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. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.</li> <li>• <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may 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.</li> <li>• <i>member</i> - Nodes that are members have successfully joined the cluster.</li> </ul>
model	string	
name	string	



Name	Type	Description
serial_number	string	
service_processor	<a href="#">service_processor</a>	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	<a href="#">version</a>	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.

#### 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.
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 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 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 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.

cluster

Complete cluster information

Name	Type	Description
_links	<a href="#">_links</a>	
configuration_backup	<a href="#">configuration_backup</a>	
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"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>
license	<a href="#">license</a>	License keys or NLF contents.
location	string	
management_interface	<a href="#">management_interface</a>	The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.
management_interfaces	array[ <a href="#">management_interfaces</a> ]	

Name	Type	Description
metric	<a href="#">metric</a>	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.
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.
statistics	<a href="#">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.
uuid	string	
version	<a href="#">version</a>	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.

#### job\_link

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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

#### Retrieving chassis information

The chassis GET API retrieves all of the chassis information in the cluster.

#### Examples

##### 1) Retrieve 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"
    }
  }
}
```

## 2) Retrieve 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",
    "_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`

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
<code>frus.type</code>	string	query	False	Filter by frus.type
<code>frus.state</code>	string	query	False	Filter by frus.state
<code>frus.id</code>	string	query	False	Filter by frus.id
<code>nodes.name</code>	string	query	False	Filter by nodes.name
<code>nodes.uuid</code>	string	query	False	Filter by nodes.uuid
<code>state</code>	string	query	False	Filter by state
<code>shelves.uuid</code>	string	query	False	Filter by shelves.uuid
<code>id</code>	string	query	False	Filter by id
<code>fields</code>	array[string]	query	False	Specify the fields to return.
<code>max_records</code>	integer	query	False	Limit the number of records returned.
<code>return_records</code>	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">chassis</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "frus": {
      "state": "ok",
      "type": "fan"
    },
    "id": "021352005981",
    "nodes": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "shelves": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uid": 7777841915827391056
    },
    "state": "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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

frus

Name	Type	Description
id	string	
state	string	
type	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

nodes

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

shelf\_reference

Shelf

Name	Type	Description
_links	<a href="#">_links</a>	
uid	string	

chassis

Name	Type	Description
frus	array[frus]	List of frus in chassis
id	string	
nodes	array[nodes]	List of nodes in 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}

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 chassis
id	string	
nodes	array[nodes]	List of nodes in chassis
shelves	array[shelf_reference]	List of shelves in chassis
state	string	

## Example response

```
{
  "frus": {
    "state": "ok",
    "type": "fan"
  },
  "id": "021352005981",
  "nodes": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "shelves": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uid": 7777841915827391056
  },
  "state": "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

frus

Name	Type	Description
id	string	
state	string	
type	string	

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

nodes

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

shelf\_reference

Shelf

Name	Type	Description
_links	<a href="#">_links</a>	
uid	string	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

#### Summary

This API is used to view and manipulate jobs. Jobs provide information about asynchronous operations. Some long-running jobs are paused or cancelled by calling PATCH. Individual operations will mention if they support PATCH on the job. Once a job transitions to a terminal state, it is deleted after a default time of 300 seconds. Attempts to GET or PATCH the job will return a 404 error code once the job has been deleted.

#### Example

The following examples show how to retrieve and update a job state

#### 1) Retrieve 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"
    }
  }
}
```

---

## 2) Update 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

Retrieves a list of recently running asynchronous jobs. Once a job transitions to a failure or success state, it is deleted after a default time of 300 seconds.

### Learn more

- [DOC /cluster/jobs](#)

## Parameters

Name	Type	In	Required	Description
state	string	query	False	Filter by state
start_time	string	query	False	Filter by start_time
end_time	string	query	False	Filter by end_time
code	integer	query	False	Filter by code
uuid	string	query	False	Filter by uuid
description	string	query	False	Filter by description
message	string	query	False	Filter by message
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	
records	array[ <a href="#">job</a> ]	

## 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": "queued",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

job

Name	Type	Description
_links	<a href="#">_links</a>	
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.
uuid	string	

error\_arguments



Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}`

Retrieves the details of a specific asynchronous job. Once a job transitions to a failure or success state, it is deleted after a default time of 300 seconds.

#### Learn more

- [DOC /cluster/jobs](#)

#### 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
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
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.
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": "queued",
  "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	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}`

Updates the state of a specific asynchronous job.

### Learn more

- [DOC /cluster/jobs](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Job UUID
action	string	query	False	<p>Request a job to pause, resume, or cancel.</p> <p>Not all jobs support these actions. A job can only be resumed if it is in paused state. Upon successfully requesting a job to be cancelled, the job state changes to either success or failure.</p> <p> Not all jobs support these actions. A job can only be resumed if it is in paused state. Upon successfully requesting a job to be cancelled, the job state changes to either success or failure.</p> <ul style="list-style-type: none"> <li>enum: ["pause", "resume", "cancel"]</li> </ul>

## 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.
458783	This job does not support pause.
458784	This job does not support cancel.

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[ <a href="#">error_arguments</a> ]	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. New features can be enabled 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 which 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
- cluster
- node

A package licensed under 'site' or 'cluster' permits the feature to be used by any node that is a member of the cluster.

A package licensed under 'node' scope permits the authorized node to use the feature. Within a cluster, if you haven't supplied every node with a valid license, the package state will indicate 'noncompliant'. A license must be purchased 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'. A new license will need to be purchased 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 can not be reached, or is not known to the cluster, the individual license may indicate an 'unknown' state.

## Licensing keys

A license is issued in one of the following two formats:

- 26-character key
- NetApp License File (NLF)

The following is an example of a 26-character key:

```
AMEPOSOIKLKGEEEEEDGNDEKSJDE
```



The following is an example of a 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"
}
```

Either format can be submitted, via this API, 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
GET /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
GET /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"
    }
  }
}
```

```
}  
}  
}
```

## Installing a NLF license

This example installs a single license in the NLF format.



You must escape all double quotes and backslash characters, of the JSON license, before it can be placed in the POST request.

```
# API  
POST /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 26-character key

This example installs a single 26-character key formatted license.

```
# API  
POST /cluster/licensing/licenses/  
  
# JSON Body  
{  
  "keys" : [ "AAAAAAAAAAAAAAAAAAAAAAAAAAAA" ]  
}  
  
# Response  
201 Created
```

## Installing multiple licenses with one API call

This example demonstrates how multiple keys can be provided to install multiple features in a single API call.

```
# API
POST /cluster/licensing/licenses/

# JSON Body
{
  "keys" : [ "AAAAAAAAAAAAAAAAAAAAAAAAAAAA",
            "BBBBBBBBBBBBBBBBBBBBBBBBBBBB" ]
}

# Response
201 Created
```

### Retrieve license packages

GET /cluster/licensing/licenses

Retrieves a collection of license packages.

### Related ONTAP commands

- `system license show-status`
- `system license show`

### Learn more

- [DOC /cluster/licensing/licenses](#)

### Parameters

Name	Type	In	Required	Description
scope	string	query	False	Filter by scope
licenses.expiry_time	string	query	False	Filter by licenses.expiry_time
licenses.compliance.state	string	query	False	Filter by licenses.compliance.state
licenses.capacity.used_size	integer	query	False	Filter by licenses.capacity.used_size

Name	Type	In	Required	Description
licenses.capacity.maximum_size	integer	query	False	Filter by licenses.capacity.maximum_size
licenses.serial_number	string	query	False	Filter by licenses.serial_number
licenses.active	boolean	query	False	Filter by licenses.active
licenses.owner	string	query	False	Filter by licenses.owner
licenses.evaluation	boolean	query	False	Filter by licenses.evaluation
licenses.start_time	string	query	False	Filter by licenses.start_time
state	string	query	False	Filter by state
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	The default is true for GET calls. When set to false, only the number of records is returned.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">records</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "keys": {
    },
    "licenses": {
      "capacity": {
        "maximum_size": 0,
        "used_size": 0
      },
      "compliance": {
        "state": "compliant"
      },
      "expiry_time": "2019-03-02 19:00:00 UTC",
      "owner": "cluster1",
      "serial_number": "123456789",
      "start_time": "2019-02-02 19:00:00 UTC"
    },
    "name": "NFS",
    "scope": "not_available",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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	<a href="#">capacity</a>	
compliance	<a href="#">compliance</a>	
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	<a href="#">_links</a>	
keys	array[string]	
licenses	array[ <a href="#">licenses</a> ]	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[ <a href="#">error_arguments</a> ]	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

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`

#### Learn more

- [DOC /cluster/licensing/licenses](#)

#### Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned.

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>keys</code>	array[string]	
<code>licenses</code>	array[ <a href="#">licenses</a> ]	Installed licenses of the package.
<code>name</code>	string	Name of the license.
<code>scope</code>	string	Scope of the license.
<code>state</code>	string	Summary state of package based on all installed licenses.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "keys": {
  },
  "licenses": {
    "capacity": {
      "maximum_size": 0,
      "used_size": 0
    },
    "compliance": {
      "state": "compliant"
    },
    "expiry_time": "2019-03-02 19:00:00 UTC",
    "owner": "cluster1",
    "serial_number": "123456789",
    "start_time": "2019-02-02 19:00:00 UTC"
  },
  "name": "NFS",
  "scope": "not_available",
  "state": "compliant"
}
```

## Response

Status: 201, Created

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">records</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "keys": {
    },
    "licenses": {
      "capacity": {
        "maximum_size": 0,
        "used_size": 0
      },
      "compliance": {
        "state": "compliant"
      },
      "expiry_time": "2019-03-02 19:00:00 UTC",
      "owner": "cluster1",
      "serial_number": "123456789",
      "start_time": "2019-02-02 19:00:00 UTC"
    },
    "name": "NFS",
    "scope": "not_available",
    "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
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
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
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

Name	Type	Description
self	<a href="#">href</a>	

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	<a href="#">capacity</a>	
compliance	<a href="#">compliance</a>	
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.

#### license\_package

Name	Type	Description
_links	<a href="#">_links</a>	
keys	array[string]	
licenses	array[ <a href="#">licenses</a> ]	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.

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### records

Name	Type	Description
_links	<a href="#">_links</a>	
keys	array[string]	
licenses	array[ <a href="#">licenses</a> ]	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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage a cluster license package

### Cluster licensing licenses name endpoint overview

#### Overview

Manages a specific instance of a license package.

#### Examples

##### Retrieving information for a specific license package

This example shows how to retrieve information about the specific feature package `fabricpool`.

```
# API
GET /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
DELETE /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
DELETE /cluster/licensing/licenses/nfs/?serial_number=*

# JSON Body
{}

# Response
200 OK
```

### Delete a license

```
DELETE /cluster/licensing/licenses/{name}
```

Deletes a license.

### Related ONTAP commands

- `system license delete`

### Learn more

- [DOC /cluster/licensing/licenses/{name}](#)

### Parameters

Name	Type	In	Required	Description
name	string	path	True	
serial_number	string	query	True	

### Response

```
Status: 200, Ok
```

### Error

```
Status: Default
```

### ONTAP Error Response Codes

Error Code	Description
1115144	Cloud licenses cannot be deleted

Error Code	Description
1115178	A tier license, that is still in use, cannot be deleted
1115406	Capacity pool licenses cannot be deleted
66846823	A flexcache licenses, that is still in use, cannot be deleted
1115213	License is still in use and cannot be removed
1115137	Cluster license requires a base license to be installed
525028	Error during volume limit check, cannot remove license
525029	Current volume use will exceed limits if license is 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[ <a href="#">error_arguments</a> ]	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}

Retrieves a specific license package.

### Related ONTAP commands

- `system license show`
- `system license show-status`

### Learn more

- [DOC /cluster/licensing/licenses/{name}](#)

### Parameters

Name	Type	In	Required	Description
name	string	path	True	Name of the license package.
scope	string	query	False	Filter by scope

Name	Type	In	Required	Description
licenses.expiry_time	string	query	False	Filter by licenses.expiry_time
licenses.compliance.state	string	query	False	Filter by licenses.compliance.state
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.serial_number	string	query	False	Filter by licenses.serial_number
licenses.active	boolean	query	False	Filter by licenses.active
licenses.owner	string	query	False	Filter by licenses.owner
licenses.evaluation	boolean	query	False	Filter by licenses.evaluation
licenses.start_time	string	query	False	Filter by licenses.start_time
state	string	query	False	Filter by state
name	string	query	False	Filter by name
fields	array[string]	query	False	Specify the fields to return.

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	



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 response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "keys": {
  },
  "licenses": {
    "capacity": {
      "maximum_size": 0,
      "used_size": 0
    },
    "compliance": {
      "state": "compliant"
    },
    "expiry_time": "2019-03-02 19:00:00 UTC",
    "owner": "cluster1",
    "serial_number": "123456789",
    "start_time": "2019-02-02 19:00:00 UTC"
  },
  "name": "NFS",
  "scope": "not_available",
  "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	<a href="#">href</a>	

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	<a href="#">capacity</a>	
compliance	<a href="#">compliance</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve historical performance metrics for the cluster

GET /cluster/metrics

Retrieves historical performance metrics for the cluster.

### Parameters

Name	Type	In	Required	Description
status	string	query	False	Filter by status
timestamp	string	query	False	Filter by timestamp
duration	string	query	False	Filter by duration
iops.total	integer	query	False	Filter by iops.total

Name	Type	In	Required	Description
iops.other	integer	query	False	Filter by iops.other
iops.read	integer	query	False	Filter by iops.read
iops.write	integer	query	False	Filter by iops.write
latency.total	integer	query	False	Filter by latency.total
latency.other	integer	query	False	Filter by latency.other
latency.read	integer	query	False	Filter by latency.read
latency.write	integer	query	False	Filter by latency.write
throughput.total	integer	query	False	Filter by throughput.total
throughput.other	integer	query	False	Filter by throughput.other
throughput.read	integer	query	False	Filter by throughput.read
throughput.write	integer	query	False	Filter by throughput.write
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"> <li>• Default value: 1</li> </ul>
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.
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
The default is true for GET calls. When set to false, only the number of records is returned.  • Default value: 1	interval	string	query	False

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">records</a> ]	

## 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": "2017-01-25 11:20:13 UTC"
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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
<a href="#">_links</a>	<a href="#">_links</a>	

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 does 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage cluster nodes

### Cluster nodes endpoint overview

#### Overview

This API is used to add nodes to a cluster, update node-specific configurations, and retrieve the current node configuration details.

#### Adding a node to a cluster

A node can be added to a cluster by issuing a POST `/cluster/nodes` request to a node currently in the cluster. All nodes must be at the same version to use this API. Mixed version joins are not supported in this release. Properties can be provided 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. The `/cluster/jobs` APIs can be used 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

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 setup additional cluster-wide configuration:

- name
- location
- records

### Network interface fields

Each node can have a node-specific configuration set in POST /cluster/nodes. If a field is provided in the body of a node, it must be provided for all nodes in the POST body.

The node management interface can be provided for each node if all node management interfaces in the cluster use the same netmask. If the node management interfaces use different netmasks, then configuration of the node management interfaces should be done using the /network/ip/interfaces API.

### The records field

Multiple nodes can be added to the cluster in one request by providing 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, no records field is needed. See 'Example usecases' for an example of how to use the records field.

### Modifying node configurations

The following fields can be used to modify a node configuration:

- name
- location

### Examples

The following examples show how to shutdown/reboot a node and how to update a node configuration.

### Adding a single node with a minimal configuration

```
# Body
body =
{
  "cluster_interface": {
    "ip": {
      "address": "1.1.1.1"
    }
  }
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster/nodes" -d body
```

## Adding multiple nodes in the same request

```
# Body
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" -d body
```

## Modifying a cluster-wide configuration

```
# Body
body =
{
  "name": "renamedNode",
  "location": "newLocation"
}

# Request
curl -X PATCH "https://<mgmt-ip>/api/cluster/nodes" -d body
```

## Shutting down a node

```
curl -X PATCH "https://<mgmt-ip>/api/cluster/nodes/{uuid}?action=shutdown"
```

## Retrieve nodes in a cluster

GET /cluster/nodes

Retrieves the nodes in the cluster.

### Learn more

- [DOC /cluster/nodes](#)

### Parameters

Name	Type	In	Required	Description
ha.auto_giveback	boolean	query	False	Filter by ha.auto_giveback
ha.enabled	boolean	query	False	Filter by ha.enabled
ha.partners.name	string	query	False	Filter by ha.partners.name
ha.partners.uuid	string	query	False	Filter by ha.partners.uuid
uptime	integer	query	False	Filter by uptime
date	string	query	False	Filter by date
membership	string	query	False	Filter by membership
serial_number	string	query	False	Filter by serial_number
controller.over_temperature	string	query	False	Filter by controller.over_temperature

Name	Type	In	Required	Description
controller.flash_cache.serial_number	string	query	False	Filter by controller.flash_cache.serial_number
controller.flash_cache.hardware_revision	string	query	False	Filter by controller.flash_cache.hardware_revision
controller.flash_cache.capacity	integer	query	False	Filter by controller.flash_cache.capacity
controller.flash_cache.model	string	query	False	Filter by controller.flash_cache.model
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.firmware_version	string	query	False	Filter by controller.flash_cache.firmware_version
controller.flash_cache.part_number	string	query	False	Filter by controller.flash_cache.part_number
controller.frus.id	integer	query	False	Filter by controller.frus.id
controller.frus.type	string	query	False	Filter by controller.frus.type
controller.frus.state	string	query	False	Filter by controller.frus.state
location	string	query	False	Filter by location
model	string	query	False	Filter by model



Name	Type	In	Required	Description
management_interfaces.ip.address	string	query	False	Filter by management_interfaces.ip.address
management_interfaces.uuid	string	query	False	Filter by management_interfaces.uuid
management_interfaces.name	string	query	False	Filter by management_interfaces.name
service_processor.link_status	string	query	False	Filter by service_processor.link_status
service_processor.state	string	query	False	Filter by service_processor.state
service_processor.firmware_version	string	query	False	Filter by service_processor.firmware_version
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.netmask	string	query	False	Filter by service_processor.ipv4_interface.netmask
service_processor.ipv4_interface.gateway	string	query	False	Filter by service_processor.ipv4_interface.gateway
service_processor.mac_address	string	query	False	Filter by service_processor.mac_address

Name	Type	In	Required	Description
service_processor.ip v6_interface.address	string	query	False	Filter by service_processor.ip v6_interface.address
service_processor.ip v6_interface.netmask	string	query	False	Filter by service_processor.ip v6_interface.netmask
service_processor.ip v6_interface.gateway	string	query	False	Filter by service_processor.ip v6_interface.gateway
name	string	query	False	Filter by name
version.minor	integer	query	False	Filter by version.minor
version.full	string	query	False	Filter by version.full
version.major	integer	query	False	Filter by version.major
version.generation	integer	query	False	Filter by version.generation
cluster_interfaces.ip. address	string	query	False	Filter by cluster_interfaces.ip. address
cluster_interfaces.uu id	string	query	False	Filter by cluster_interfaces.uu id
cluster_interfaces.na me	string	query	False	Filter by cluster_interfaces.na me
uuid	string	query	False	Filter by 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_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	
records	array[ <a href="#">records</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "cluster_interface": {
      "ip": {
        "address": "10.10.10.7"
      }
    },
    "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": 102400000000,
        "firmware_version": "NA05",
        "hardware_revision": "A1",
        "model": "X1970A",
        "part_number": "119-00207",
        "serial_number": "A22P5061550000187",
        "slot": "6-1",
        "state": "ok"
      },
      "frus": {
```

```
    "id": 0,  
    "state": "ok",  
    "type": "fan"  
  },  
  "over_temperature": "over"  
},  
"date": "2017-01-25 11:20:13 +0400",  
"ha": {  
  "partners": {  
    "_links": {  
      "self": {  
        "href": "/api/resourcelink"  
      }  
    },  
    "name": "node1",  
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"  
  }  
},  
"location": "rack 2 row 5",  
"management_interface": {  
  "ip": {  
    "address": "10.10.10.7"  
  }  
},  
"management_interfaces": {  
  "_links": {  
    "self": {  
      "href": "/api/resourcelink"  
    }  
  },  
  "ip": {  
    "address": "10.10.10.7"  
  },  
  "name": "lif1",  
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"  
},  
"membership": "available",  
"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": "up",
    "mac_address": "string",
    "state": "online"
  },
  "uptime": 300536,
  "uuid": "4ea7a442-86d1-11e0-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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster\_interfaces

Network interface



Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	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[ <a href="#">flash_cache</a> ]	A list of Flash-Cache devices. Only returned when requested by name.
frus	array[ <a href="#">frus</a> ]	A list of frus in 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.

#### partners

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	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.
partners	array[ <a href="#">partners</a> ]	The nodes in this node's High Availability (HA) group.

#### management\_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

#### management\_interfaces

Network interface

Name	Type	Description
_links	<a href="#">_links</a>	

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

#### 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.

Name	Type	Description
firmware_version	string	The version of firmware installed.
ipv4_interface	<a href="#">ipv4_interface</a>	Object to setup an interface along with its default router.
ipv6_interface	<a href="#">ipv6_interface</a>	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

### 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.

### records

#### Complete node information

Name	Type	Description
_links	<a href="#">_links</a>	
cluster_interface	<a href="#">cluster_interface</a>	The cluster network IP address of the node to be added.
cluster_interfaces	array[ <a href="#">cluster_interfaces</a> ]	
controller	<a href="#">controller</a>	Controller information
date	string	Specifies the ISO-8601 format date and time on the node.
ha	<a href="#">ha</a>	

Name	Type	Description
location	string	
management_interface	<a href="#">management_interface</a>	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[ <a href="#">management_interfaces</a> ]	
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> <li>• <i>available</i> - If a node is available, this means it 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. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.</li> <li>• <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may 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.</li> <li>• <i>member</i> - Nodes that are members have successfully joined the cluster.</li> </ul>
model	string	
name	string	
serial_number	string	
service_processor	<a href="#">service_processor</a>	

Name	Type	Description
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	<a href="#">version</a>	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[ <a href="#">error_arguments</a> ]	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`

Adds a node or nodes to the cluster

#### Required properties

- `cluster_interface.ip.address`

#### Learn more

- [DOC /cluster/nodes](#)

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
cluster_interface	<a href="#">cluster_interface</a>	The cluster network IP address of the node to be added.
cluster_interfaces	array[ <a href="#">cluster_interfaces</a> ]	
controller	<a href="#">controller</a>	Controller information
date	string	Specifies the ISO-8601 format date and time on the node.
ha	<a href="#">ha</a>	
location	string	
management_interface	<a href="#">management_interface</a>	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[ <a href="#">management_interfaces</a> ]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> <li>• <i>available</i> - If a node is available, this means it 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. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.</li> <li>• <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may 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.</li> <li>• <i>member</i> - Nodes that are members have successfully joined the cluster.</li> </ul>
model	string	
name	string	
serial_number	string	
service_processor	<a href="#">service_processor</a>	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	<a href="#">version</a>	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 request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "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": 102400000000,
      "firmware_version": "NA05",
      "hardware_revision": "A1",
      "model": "X1970A",
      "part_number": "119-00207",
      "serial_number": "A22P5061550000187",
      "slot": "6-1",
      "state": "ok"
    },
    "frus": {
      "id": 0,
      "state": "ok",
      "type": "fan"
    },
    "over_temperature": "over"
  },
  "date": "2017-01-25 11:20:13 +0400",
  "ha": {
    "partners": {
```

```
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
},
"location": "rack 2 row 5",
"management_interface": {
  "ip": {
    "address": "10.10.10.7"
  }
},
"management_interfaces": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip": {
    "address": "10.10.10.7"
  },
  "name": "lif1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"membership": "available",
"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": "up",
  "mac_address": "string",
  "state": "online"
}
```

```

    },
    "uptime": 300536,
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
    "version": {
      "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
      "generation": 9,
      "major": 4,
      "minor": 0
    }
  }
}

```

### Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

### Example response

```

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

```

### Error

Status: Default

### ONTAP Error Response Codes

Error Code	Description
131727360	A node cannot be added to the cluster. This is a generic code, see response message for details.

Error Code	Description
262245	The value provided was invalid.
1179817	The IP address, netmask, and gateway must all be provided for cluster management interface.
1179813	Fields set for one node must be set for all nodes.
1179818	The IP address and gateway must be of the same family.
1179821	An IP address and netmask conflicts with an existing entry.
1179795	A node being added is already 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster\_interfaces

Network interface

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
name	string	The name of the interface.

Name	Type	Description
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]	A list of frus in 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.

## partners

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	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.
partners	array[ <a href="#">partners</a> ]	The nodes in this node's High Availability (HA) group.

## management\_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

## management\_interfaces

### Network interface

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

## 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	<a href="#">ipv4_interface</a>	Object to setup an interface along with its default router.
ipv6_interface	<a href="#">ipv6_interface</a>	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	



Name	Type	Description
state	string	

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.

node

Complete node information

Name	Type	Description
_links	<a href="#">_links</a>	
cluster_interface	<a href="#">cluster_interface</a>	The cluster network IP address of the node to be added.
cluster_interfaces	array[ <a href="#">cluster_interfaces</a> ]	
controller	<a href="#">controller</a>	Controller information
date	string	Specifies the ISO-8601 format date and time on the node.
ha	<a href="#">ha</a>	
location	string	
management_interface	<a href="#">management_interface</a>	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[ <a href="#">management_interfaces</a> ]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> <li>• <i>available</i> - If a node is available, this means it 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. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.</li> <li>• <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may 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.</li> <li>• <i>member</i> - Nodes that are members have successfully joined the cluster.</li> </ul>
model	string	
name	string	
serial_number	string	
service_processor	<a href="#">service_processor</a>	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	<a href="#">version</a>	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.

## job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}`

Retrieves information for the node.

### Learn more

- [DOC /cluster/nodes](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	• format: uuid
fields	array[string]	query	False	Specify the fields to return.

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
cluster_interface	<a href="#">cluster_interface</a>	The cluster network IP address of the node to be added.
cluster_interfaces	array[ <a href="#">cluster_interfaces</a> ]	
controller	<a href="#">controller</a>	Controller information
date	string	Specifies the ISO-8601 format date and time on the node.
ha	<a href="#">ha</a>	
location	string	
management_interface	<a href="#">management_interface</a>	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[ <a href="#">management_interfaces</a> ]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> <li>• <i>available</i> - If a node is available, this means it 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. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.</li> <li>• <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may 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.</li> <li>• <i>member</i> - Nodes that are members have successfully joined the cluster.</li> </ul>
model	string	
name	string	
serial_number	string	
service_processor	<a href="#">service_processor</a>	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	<a href="#">version</a>	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"
    }
  },
  "cluster_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "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": 102400000000,
      "firmware_version": "NA05",
      "hardware_revision": "A1",
      "model": "X1970A",
      "part_number": "119-00207",
      "serial_number": "A22P5061550000187",
      "slot": "6-1",
      "state": "ok"
    },
    "frus": {
      "id": 0,
      "state": "ok",
      "type": "fan"
    },
    "over_temperature": "over"
  },
  "date": "2017-01-25 11:20:13 +0400",
  "ha": {
    "partners": {
```

```
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
},
"location": "rack 2 row 5",
"management_interface": {
  "ip": {
    "address": "10.10.10.7"
  }
},
"management_interfaces": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip": {
    "address": "10.10.10.7"
  },
  "name": "lif1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"membership": "available",
"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": "up",
  "mac_address": "string",
  "state": "online"
}
```

```
},
"uptime": 300536,
"uuid": "4ea7a442-86d1-11e0-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	<a href="#">href</a>	

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	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster\_interfaces

Network interface

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
name	string	The name of the interface.

Name	Type	Description
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[ <a href="#">flash_cache</a> ]	A list of Flash-Cache devices. Only returned when requested by name.
frus	array[ <a href="#">frus</a> ]	A list of frus in 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.

## partners

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	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.
partners	array[ <a href="#">partners</a> ]	The nodes in this node's High Availability (HA) group.

## management\_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

## management\_interfaces

### Network interface

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

## 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	<a href="#">ipv4_interface</a>	Object to setup an interface along with its default router.
ipv6_interface	<a href="#">ipv6_interface</a>	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	

Name	Type	Description
state	string	

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[ <a href="#">error_arguments</a> ]	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}

Updates the node information or performs shutdown/reboot actions on a node.

## Learn more

- [DOC /cluster/nodes](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	<ul style="list-style-type: none"><li>• format: uuid</li></ul>
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.</p> <ul style="list-style-type: none"><li>• enum: ["shutdown", "reboot"]</li></ul>
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"><li>• Default value:</li></ul>

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
cluster_interface	<a href="#">cluster_interface</a>	The cluster network IP address of the node to be added.
cluster_interfaces	array[ <a href="#">cluster_interfaces</a> ]	
controller	<a href="#">controller</a>	Controller information
date	string	Specifies the ISO-8601 format date and time on the node.
ha	<a href="#">ha</a>	
location	string	
management_interface	<a href="#">management_interface</a>	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[ <a href="#">management_interfaces</a> ]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> <li>• <i>available</i> - If a node is available, this means it 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. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.</li> <li>• <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may 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.</li> <li>• <i>member</i> - Nodes that are members have successfully joined the cluster.</li> </ul>
model	string	
name	string	
serial_number	string	
service_processor	<a href="#">service_processor</a>	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	<a href="#">version</a>	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 request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "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": 102400000000,
      "firmware_version": "NA05",
      "hardware_revision": "A1",
      "model": "X1970A",
      "part_number": "119-00207",
      "serial_number": "A22P5061550000187",
      "slot": "6-1",
      "state": "ok"
    },
    "frus": {
      "id": 0,
      "state": "ok",
      "type": "fan"
    },
    "over_temperature": "over"
  },
  "date": "2017-01-25 11:20:13 +0400",
  "ha": {
    "partners": {
```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"location": "rack 2 row 5",
"management_interface": {
  "ip": {
    "address": "10.10.10.7"
  }
},
"management_interfaces": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip": {
    "address": "10.10.10.7"
  },
  "name": "lif1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"membership": "available",
"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": "up",
  "mac_address": "string",
  "state": "online"
}

```

```

},
"uptime": 300536,
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
"version": {
  "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
  "generation": 9,
  "major": 4,
  "minor": 0
}
}

```

### 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
852046	HA partner node
65562	Internal RPC error

Error Code	Description
852115	The reboot/shutdown is prevented because LIFs cannot be moved away from the node
9240606	The reboot/shutdown is prevented due to quorum warnings.

Name	Type	Description
error	<a href="#">error</a>	

### 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	<a href="#">href</a>	

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	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster\_interfaces

Network interface

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
name	string	The name of the interface.

Name	Type	Description
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]	A list of frus in 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.

## partners

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	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.
partners	array[ <a href="#">partners</a> ]	The nodes in this node's High Availability (HA) group.

## management\_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	<a href="#">node_setup_ip</a>	The IP configuration for cluster setup.

## management\_interfaces

### Network interface

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

## 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	<a href="#">ipv4_interface</a>	Object to setup an interface along with its default router.
ipv6_interface	<a href="#">ipv6_interface</a>	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	



Name	Type	Description
state	string	

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.

node

Complete node information

Name	Type	Description
_links	<a href="#">_links</a>	
cluster_interface	<a href="#">cluster_interface</a>	The cluster network IP address of the node to be added.
cluster_interfaces	array[ <a href="#">cluster_interfaces</a> ]	
controller	<a href="#">controller</a>	Controller information
date	string	Specifies the ISO-8601 format date and time on the node.
ha	<a href="#">ha</a>	
location	string	
management_interface	<a href="#">management_interface</a>	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[ <a href="#">management_interfaces</a> ]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> <li>• <i>available</i> - If a node is available, this means it 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. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created.</li> <li>• <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may 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.</li> <li>• <i>member</i> - Nodes that are members have successfully joined the cluster.</li> </ul>
model	string	
name	string	
serial_number	string	
service_processor	<a href="#">service_processor</a>	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	<a href="#">version</a>	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.

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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

#### Cluster peer operations

Cluster peering allows administrators of ONTAP systems to establish relationships between two or more independent clusters. Once a relationship exists between two clusters, they may then exchange user data, configuration information and coordinate operations. The `/cluster/peers` endpoint supports operations such as create, get, modify and delete using GET, PATCH and POST HTTP requests.

#### Creating a cluster peer

A new cluster peer relationship can be set up 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 a HTTP status code, code 201, along with the details of the created peer such as peer UUID, name, 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 request 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 = {}

# Create with a peer address and a passphrase
body =
{
  "remote":
  {
    "ip_addresses":["1.2.3.4"]
  }
}

# Create with a peer name and a generated passphrase that is true
body =
{
  "name":"cp_xyz123",
  "authentication":
  {
    "generate_passphrase":true
  }
}

# Create with a name, a peer address, and a passphrase
body =
{
  "name":"cp_xyz123",
  "remote":
  {
    "ip_addresses": ["1.2.3.4"]
  },
  "authentication":
  {
    "passphrase":"xyz12345"
  }
}

# Create with a proposed encryption protocol
body =
{
  "encryption":
  {
    "proposed":"tls-psk"
  }
}
```

## Creating 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 relationship would be established if there is an error preventing the LIFs from being created.

Local interfaces, once created, should not be specified 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 netmask 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

It might happen that when creating LIFs the network route discovery mechanism could 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

```
curl -X POST "https://<mgmt-ip>/api/cluster/peers" -d body
```

where "<mgmt-ip>" is replaced by the IP address of the cluster management LIF, and "body" is replaced by the JSON body of the POST, containing the fields for the new peering relationship and local LIFs.</mgmt-ip>

### Example POST body

To create 4 intercluster LIFs on a 4-node cluster before creating a cluster peer relationship:

```

{
  "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"
}

```

### Retrieve a cluster peer

Peers in a cluster can be retrieved 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}`.

### Overview of fields used for retrieving a cluster peer

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.

### 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

A cluster peer relationship can be updated 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.

## Fields overview

This sections highlights 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



```
# Update with an empty body
body = {}

# Update the proposed encryption protocol from tls-psk to none
body =
{
  "authentication":
  {
    "passphrase": "xyz12345",
    "in_use": "ok"
  },
  "encryption":
  {
    "proposed": "none"
  }
}

# Update the passphrase
body =
{
  "authentication":
  {
    "passphrase": "xyz12345",
    "in_use": "ok"
  }
}

# Set an auto-generated passphrase
body =
{
  "authentication":
  {
    "generate_passphrase": true,
    "in_use": "ok"
  }
}

# Update remote IP addresses
body =
{
  "remote":
  {
    "ip_addresses": ["10.224.65.30"]
  }
}
```

## Sample requests

```
# Set a passphrase
curl -X PATCH 'https://<mgmt-ip>/api/cluster/peers/73123071-d0b9-11e8-
a686-005056a7179a' -d
'{"authentication":{"passphrase":"xyz12345","in_use":"ok"}}'

# Update 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

This interface allows you to delete a cluster peer using the HTTP DELETE request.

### Required fields

All delete operations must be performed on a valid peer UUID. Deleting an invalid peer returns 'HTTP 404' indicating an error.

### Optional fields

The DELETE operation has no optional fields.

### Request format

DELETE "https://<mgmt-ip>/api/cluster/peers/{uuid}"</mgmt-ip>

### Examples

The request -

```
curl -X DELETE "https://<mgmt-ip>/api/cluster/peers/8becc0d4-c12c-11e8-
9ceb-005056bbd143"
```

deletes a peer with peer UUID '8becc0d4-c12c-11e8-9ceb-005056bbd143'

### Retrieve cluster peers

GET /cluster/peers

Retrieve the collection of cluster peers.

### Learn more

- [DOC /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	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">cluster_peer</a> ]	

## 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": "ok",
      "state": "ok"
    },
    "encryption": {
      "proposed": "none",
      "state": "none"
    },
    "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"
    }
  }
}
```

```
},
"local_network": {
  "broadcast_domain": "bd1",
  "gateway": "10.1.1.1",
  "interfaces": {
    "ip_address": "10.10.10.7"
  },
  "netmask": "255.255.0.0"
},
"name": "cluster2",
"remote": {
  "ip_addresses": {
  },
  "name": "cluster2",
  "serial_number": "4048820-60-9"
},
"status": {
  "state": "available",
  "update_time": "2017-01-25 11:20:13 UTC"
},
"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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
interfaces	<a href="#">href</a>	
self	<a href="#">href</a>	

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	<a href="#">href</a>	

initial\_allowed\_svms

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
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[ <a href="#">interfaces</a> ]	
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.

## cluster\_peer

Name	Type	Description
_links	<a href="#">_links</a>	
authentication	<a href="#">authentication</a>	
encryption	<a href="#">encryption</a>	
initial_allowed_svms	array[ <a href="#">initial_allowed_svms</a> ]	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	<a href="#">ipspace</a>	The IPspace of the local intercluster LIFs
local_network	<a href="#">local_network</a>	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.

Name	Type	Description
remote	<a href="#">remote</a>	
status	<a href="#">status</a>	
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create a peering relationship

POST `/cluster/peers`

Creates a peering relationship and, optionally, the IP interfaces it will use. There are two ways to create a peering relationship.

### Provide remote IP

Here the user provides the remote IP address. Creating a new cluster peer relationship with a specific remote cluster requires at least one remote intercluster IP address from that cluster.

### 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.

- Either set `generate_passphrase` to true or provide a passphrase in the body of the request; only one of them is required.

### Optional properties

The following fields are optional for a POST on `/cluster/peer/`. All fields must follow the structure in the cluster peer API definition.

- `name` - Name of the peering relationship.
- `passphrase` - User generated passphrase for use in authentication.
- `generate_passphrase` (true/false) - When this option is 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` - 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.
- `local_network` - fields to create a local intercluster LIF. See section on "Creating local intercluster lifs".
- `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" should always be present and "T" should 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.

### Do not provide remote IP

This method is used when the remote IP address is not provided. This method is used when the filer is ready to accept peering request from foreign clusters.

### Required properties

- `generate_passphrase` (true) - This option must be set to true. ONTAP automatically generates a passphrase to authenticate cluster peers. Either set `generate_passphrase` to true or provide a passphrase in the body of the request; only one of them is required.

### Optional properties

The following fields are optional for a POST on `/cluster/peer/`. All fields must follow the structure in the cluster peer API definition.

- `name` - Name of the remote peer.
- `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. This list can be modified until the remote cluster accepts this cluster peering relationship.
- `local_network` - Fields to create a local intercluster LIF. See section on "Creating local intercluster lifs".
- `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" should always be present and "T" should 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.

#### 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)
- 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 initial allowed vservers peer names or UUIDs. (OK)

#### Learn more

- [DOC /cluster/peers](#)

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>authentication</code>	<a href="#">authentication</a>	
<code>encryption</code>	<a href="#">encryption</a>	
<code>initial_allowed_svms</code>	<code>array[<a href="#">initial_allowed_svms</a>]</code>	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.
<code>ipspace</code>	<a href="#">ipspace</a>	The IPspace of the local intercluster LIFs
<code>local_network</code>	<a href="#">local_network</a>	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.
remote	<a href="#">remote</a>	
status	<a href="#">status</a>	
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

```
{
  "_links": {
    "interfaces": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "in_use": "ok",
    "state": "ok"
  },
  "encryption": {
    "proposed": "none",
    "state": "none"
  },
  "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"
  },
  "local_network": {
    "broadcast_domain": "bd1",
    "gateway": "10.1.1.1",
    "interfaces": {
      "ip_address": "10.10.10.7"
    },
    "netmask": "255.255.0.0"
  },
}
```

```

"name": "cluster2",
"remote": {
  "ip_addresses": {
  },
  "name": "cluster2",
  "serial_number": "4048820-60-9"
},
"status": {
  "state": "available",
  "update_time": "2017-01-25 11:20:13 UTC"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}

```

### Response

Status: 201, Created

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
authentication	<a href="#">authentication</a>	
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 may be autogenerated for anonymous cluster peer offers.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resource/link"
    }
  },
  "authentication": {
    "expiry_time": "2017-01-25 11:20:13 UTC"
  },
  "ip_address": "10.10.10.7",
  "name": "cluster2"
}
```

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
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.
4656077	Specify either remote IP addresses or generate_passphrase.
4656075	Cannot specify encryption: this operation requires an ECV of 9.6.0 or later.
4656079	No cluster nodes were found. Check your cluster configuration.
4656085	Cannot create an intercluster LIF with an empty list of local IP addresses.



Error Code	Description
4656087	The number of local intercluster IP addresses must be less than or equal to the number of available nodes.
4656086	Creating an intercluster LIF requires a broadcast domain that is in use within the IPspace.
4653365	IPspaces are unavailable with cluster peering: {ipspace}.
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 network mask or a network mask length.
4656081	Creating an intercluster LIF requires a list of local IP addresses.
1966366	The System SVM of the cluster IPspace hosts Cluster LIFs only.
4656096	Creating an intercluster LIF requires an IPv4 or IPv6 address of the default router.

Name	Type	Description
error	<a href="#">error</a>	

### 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	<a href="#">href</a>	
self	<a href="#">href</a>	

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	<a href="#">href</a>	

initial\_allowed\_svms

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### ipospace

The IPspace of the local intercluster LIFs

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">interfaces</a> ]	
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.

cluster\_peer

Name	Type	Description
_links	<a href="#">_links</a>	
authentication	<a href="#">authentication</a>	
encryption	<a href="#">encryption</a>	
initial_allowed_svms	array[ <a href="#">initial_allowed_svms</a> ]	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	<a href="#">ipspace</a>	The IPspace of the local intercluster LIFs
local_network	<a href="#">local_network</a>	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.
remote	<a href="#">remote</a>	
status	<a href="#">status</a>	

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.

#### 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[ <a href="#">error_arguments</a> ]	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}

Deletes a cluster peer.

## Learn more

- [DOC /cluster/peers](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

## 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	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	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}`

Retrieves a specific cluster peer instance.

### Learn more

- [DOC /cluster/peers](#)

### 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	<a href="#">_links</a>	
authentication	<a href="#">authentication</a>	
encryption	<a href="#">encryption</a>	
initial_allowed_svms	array[ <a href="#">initial_allowed_svms</a> ]	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	<a href="#">ipspace</a>	The IPspace of the local intercluster LIFs
local_network	<a href="#">local_network</a>	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.
remote	<a href="#">remote</a>	
status	<a href="#">status</a>	
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 response

```
{
  "_links": {
    "interfaces": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "in_use": "ok",
    "state": "ok"
  },
  "encryption": {
    "proposed": "none",
    "state": "none"
  },
  "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"
  },
  "local_network": {
    "broadcast_domain": "bd1",
    "gateway": "10.1.1.1",
    "interfaces": {
      "ip_address": "10.10.10.7"
    },
    "netmask": "255.255.0.0"
  },
}
```

```
"name": "cluster2",
"remote": {
  "ip_addresses": {
  },
  "name": "cluster2",
  "serial_number": "4048820-60-9"
},
"status": {
  "state": "available",
  "update_time": "2017-01-25 11:20:13 UTC"
},
"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
interfaces	<a href="#">href</a>	
self	<a href="#">href</a>	

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	<a href="#">href</a>	

initial\_allowed\_svms

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
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[ <a href="#">interfaces</a> ]	
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.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update a cluster peer instance

PATCH /cluster/peers/{uuid}

Updates a cluster peer instance.

Learn more

- [DOC /cluster/peers](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Cluster peer relationship UUID

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
authentication	<a href="#">authentication</a>	
encryption	<a href="#">encryption</a>	
initial_allowed_svms	array[ <a href="#">initial_allowed_svms</a> ]	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	<a href="#">ipspace</a>	The IPspace of the local intercluster LIFs
local_network	<a href="#">local_network</a>	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.
remote	<a href="#">remote</a>	
status	<a href="#">status</a>	
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

```
{
  "_links": {
    "interfaces": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "in_use": "ok",
    "state": "ok"
  },
  "encryption": {
    "proposed": "none",
    "state": "none"
  },
  "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"
  },
  "local_network": {
    "broadcast_domain": "bd1",
    "gateway": "10.1.1.1",
    "interfaces": {
      "ip_address": "10.10.10.7"
    },
    "netmask": "255.255.0.0"
  },
}
```

```
"name": "cluster2",
"remote": {
  "ip_addresses": {
  },
  "name": "cluster2",
  "serial_number": "4048820-60-9"
},
"status": {
  "state": "available",
  "update_time": "2017-01-25 11:20:13 UTC"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
authentication	<a href="#">authentication</a>	
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 may be autogenerated for anonymous cluster peer offers.



## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourceLink"
    }
  },
  "authentication": {
    "expiry_time": "2017-01-25 11:20:13 UTC"
  },
  "ip_address": "10.10.10.7",
  "name": "cluster2"
}
```

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
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.
4656073	Changing the encryption state requires the refreshing of the authentication passphrase.
4656076	Cluster peer modify was attempted mismatched IPv4 and IPv6 addresses.
4656075	Cannot specify encryption: this operation requires an ECV of Data ONTAP 9.6.0 or later.
4656084	Passphrase can only be modified with an authenticated cluster peer relationship.
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.
4656081	The remote IP address list is empty.

Error Code	Description
4656092	Cluster peer modify was attempted with a host name that did not resolve to an IPv4 or IPv6 address.
4655058	Expiration time cannot be more than 7 days in the future.
4653261	Error finding IPspace.
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.

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	<a href="#">href</a>	
self	<a href="#">href</a>	

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	<a href="#">href</a>	

initial\_allowed\_svms

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### ipospace

The IPspace of the local intercluster LIFs

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">interfaces</a> ]	
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.

cluster\_peer

Name	Type	Description
_links	<a href="#">_links</a>	
authentication	<a href="#">authentication</a>	
encryption	<a href="#">encryption</a>	
initial_allowed_svms	array[ <a href="#">initial_allowed_svms</a> ]	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	<a href="#">ipspace</a>	The IPspace of the local intercluster LIFs
local_network	<a href="#">local_network</a>	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.
remote	<a href="#">remote</a>	
status	<a href="#">status</a>	

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.

#### 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[ <a href="#">error_arguments</a> ]	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

The `/cluster/schedules` API is used to view, create, and modify job schedules in a cluster.

## Retrieving a job schedule

Job schedules can be retrieved 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}`. Queries on fields can be applied 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

A job schedule is created 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. You must 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

### Example

```
# The API:
/api/cluster/schedules

# The call:
curl -X POST "https://<mgmt-ip>/api/cluster/schedules" -d body

# The response of a successful POST is empty.
Example body to create an interval schedule with a 1-week interval:
{
  "name": "test_interval_1",
  "interval": "P1W"
}
Example body to create a cron schedule that runs daily at 12:05 :
{
  "name": "test_cron_1",
  "cron":
  {
    "minutes": [ 5 ],
    "hours": [ 12 ]
  }
}
```

### Optional fields

By default, the schedule is owned by the local cluster. In a MetroCluster configuration, the partner cluster can be specified 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

Multiple schedules can be created 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: The name, cluster, and type of schedule cannot be modified. Also, you cannot modify a cron field of an interval schedule, or the interval field of a cron schedule. Queries on fields can be applied to modify all schedules that match the combined query.

## Example

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/cluster/schedules/{uuid}" -d body

# The response of a successful PATCH is empty.
Example body to modify an interval schedule with a 2-day and 5-minute
interval:
{
  "interval": "P2DT5M"
}
Example body to modify a cron schedule to run Mondays at 2:
{
  "cron":
  {
    "hours": [ 2 ],
    "weekdays": [ 1 ]
  }
}
```

## Deleting a job schedule

Job schedules can be deleted based on their UUID. Queries on fields can be applied 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

Retrieves a schedule.

## Learn more

- [DOC /cluster/schedules](#)

## Parameters

Name	Type	In	Required	Description
cron.hours	integer	query	False	Filter by cron.hours
cron.months	integer	query	False	Filter by cron.months
cron.minutes	integer	query	False	Filter by cron.minutes
cron.weekdays	integer	query	False	Filter by cron.weekdays
cron.days	integer	query	False	Filter by cron.days
cluster.name	string	query	False	Filter by cluster.name
cluster.uuid	string	query	False	Filter by cluster.uuid
interval	string	query	False	Filter by interval
uuid	string	query	False	Filter by uuid
type	string	query	False	Filter by 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	The default is true for GET calls. When set to false, only the number of records is returned.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">schedule</a> ]	

## 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": {
      },
      "hours": {
      },
      "minutes": {
      },
      "months": {
      },
      "weekdays": {
      }
    },
    "interval": "P1DT2H3M4S",
    "type": "cron",
    "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	<a href="#">error</a>	

### 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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	<a href="#">_links</a>	
cluster	<a href="#">cluster</a>	The cluster that owns the schedule. Defaults to the local cluster.
cron	<a href="#">cron</a>	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[ <a href="#">error_arguments</a> ]	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`

Create a schedule.

### Required Fields

- name - Name of the job schedule  
It is required to provide a minutes field for a cron schedule. An interval field is required for an interval schedule. You must not provide both a cron field and an interval field.

### Learn more

- [DOC /cluster/schedules](#)

### Request Body

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>cluster</code>	<code>cluster</code>	The cluster that owns the schedule. Defaults to the local cluster.
<code>cron</code>	<code>cron</code>	Details for schedules of type cron.
<code>interval</code>	string	An ISO-8601 duration formatted string.
<code>name</code>	string	Schedule Name. Required in the URL or POST body.
<code>type</code>	string	Schedule type
<code>uuid</code>	string	Job schedule UUID

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cron": {
    "days": {
    },
    "hours": {
    },
    "minutes": {
    },
    "months": {
    },
    "weekdays": {
    }
  },
  "interval": "P1DT2H3M4S",
  "type": "cron",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

## Response

Status: 201, Created

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
459760	The schedule specified is not a valid schedule.
458788	The schedule specified is not a valid schedule.

Error Code	Description
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.
459763	Schedule cannot be created locally using the remote cluster name as the owner.
460784	An error occurred creating the remote cluster version of this 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
self	<a href="#">href</a>	

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
_links	<a href="#">_links</a>	
cluster	<a href="#">cluster</a>	The cluster that owns the schedule. Defaults to the local cluster.
cron	<a href="#">cron</a>	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[ <a href="#">error_arguments</a> ]	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}
```

Deletes a schedule.

#### Learn more

- [DOC /cluster/schedules](#)

#### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

#### 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	<a href="#">error</a>	



## 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[ <a href="#">error_arguments</a> ]	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}

Retrieves a schedule.

## Learn more

- [DOC /cluster/schedules](#)

## 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	<a href="#">_links</a>	
cluster	<a href="#">cluster</a>	The cluster that owns the schedule. Defaults to the local cluster.
cron	<a href="#">cron</a>	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": {
    },
    "hours": {
    },
    "minutes": {
    },
    "months": {
    },
    "weekdays": {
    }
  },
  "interval": "P1DT2H3M4S",
  "type": "cron",
  "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	<a href="#">href</a>	

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[ <a href="#">error_arguments</a> ]	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}`

Updates a schedule. Note that you cannot modify a cron field of an interval schedule, or the interval field of a cron schedule.

### Learn more

- [DOC /cluster/schedules](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Schedule UUID

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
cluster	<a href="#">cluster</a>	The cluster that owns the schedule. Defaults to the local cluster.
cron	<a href="#">cron</a>	Details for schedules of type cron.

Name	Type	Description
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 request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cron": {
    "days": {
    },
    "hours": {
    },
    "minutes": {
    },
    "months": {
    },
    "weekdays": {
    }
  },
  "interval": "P1DT2H3M4S",
  "type": "cron",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

### Response

Status: 200, Ok

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
459760	The schedule specified is not a valid schedule.
458788	The schedule specified is not a valid schedule.
460783	As this is a MetroCluster configuration and the local cluster is waiting for switchback, changes to non-system schedules are not allowed.
459761	Schedule cannot be modified on this cluster because it is replicated from the remote 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	<a href="#">href</a>	

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
<a href="#">_links</a>	<a href="#">_links</a>	
cluster	<a href="#">cluster</a>	The cluster that owns the schedule. Defaults to the local cluster.
cron	<a href="#">cron</a>	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[ <a href="#">error_arguments</a> ]	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

ONTAP cluster software API retrieves and displays relevant information about the software profile, software packages collection, and software history collection. The API retrieves the information about all software packages present in the cluster, or specific software package.

The POST request provides the ability to download a software package from an HTTP or FTP server. The PATCH request provides the option to upgrade the cluster software version. The client can validate the package before triggering the update by selecting the `validate_only` field. Setting the `version` field triggers the installation of the package in the cluster. The client can pause, resume, or cancel any ongoing software upgrade by selecting `action`. The DELETE request can remove a specific software package present in the cluster.

### Examples

#### Retrieving software profile information

The following example shows how to retrieve software profile information. The client can check the validation results after selecting `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",
      "message": [
        {
          "code": 166,
          "message": "Use NFS hard mounts, if possible.",
          "arguments": [
            "string"
          ]
        }
      ],
      "action": [
```

```

    {
      "code": 166,
      "message": "Use NFS hard mounts, if possible.",
      "arguments": [
        "string"
      ]
    }
  ]
}
],
"version": "9.5.0",
"pending_version": "9.6.0",
"nodes": [
  {
    "node": "sti70-vsim-ucs165n",
    "version": "9.5.0"
  }
],
"metrocluster": {
  "progress_summary": "Update paused by user",
  "progress_details": "Installing Data ONTAP 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",
      "message": "Image update complete",
      "action": "",
      "start_time": "2018-05-21T09:53:04+05:30",
      "end_time": "2018-05-21T11:53:04+05:30",
      "node": {
        "name": "sti70-vsimg-ucs165n"
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/cluster/software/"
    }
  }
}
```

---

## Upgrading the software version

The following example shows how to upgrade cluster software. Setting the `version` field triggers the installation of the package. The client 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` to an ongoing upgrade. An upgrade can only be resumed if it is in the paused state.

The client can start the upgrade process at the cluster-level. There are no options available to start the upgrade for a specific node or HA pairs.

---

### 1. Validating the package and verifying the validation results

The following example shows how to validate a cluster software package. The client has to validate the package before the software upgrade. The client must set the `validate_only` field to `true` to start the validation. The client 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 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 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"
    }
  }
}
```

The client 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",
      "message": "Cluster HA is not configured in the cluster. Storage failover is not enabled on node \"node1\", \"node2\".",
      "action": "Check cluster HA configuration. Check storage failover status."
    }
  ]
}
```

```

},
{
  "update_check": "Manual checks",
  "status": "warning",
  "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": "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. The client must validate the package before the software upgrade starts. The client must set the `skip_warnings` field to `true` in order to skip any validation warnings and start the software package upgrade.



```
# 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 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 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"
    }
  }
}

```

The client 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",
      "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": "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",
      "message": "",
      "action": "",
      "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",
      "message": "",
      "action": "",
      "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/cancelling the upgrade

The following example shows how to pause an ongoing cluster software package upgrade. The client must set the `action` field to `pause`, `resume`, or `cancel` which pauses, resumes or cancels 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 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 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"
    }
  }
}
```

---

The client 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",
      "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": "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",
    "message": "Installing Data ONTAP software image.",
    "action": "",
    "start_time": "2019-01-08T02:54:36+05:30",
    "node": {
      "name": "node1"
    }
  },
  {
    "status": "in-progress",
    "message": "Installing Data ONTAP software image.",
    "action": "",
    "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. The client provides the url, username, and password 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 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 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"
    }
  }
}
```

---

### 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. The client needs to provide the package version that they 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 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 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
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-vsims-ucs165n",
    "_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

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
fields	array[string]	query	False	Specify the fields to return.
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.

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">metrocluster</a>	
nodes	array[ <a href="#">software_node_reference</a> ]	List of nodes and active versions.



Name	Type	Description
pending_version	string	Version being installed on the system. <ul style="list-style-type: none"> <li>• example: ONTAP_X_1</li> <li>• readOnly: 1</li> </ul>
state	string	Operational state of the upgrade
status_details	array[ <a href="#">software_status_details_reference</a> ]	Display status details.
update_details	array[ <a href="#">software_update_details_reference</a> ]	Display update process details.
validation_results	array[ <a href="#">software_validation_reference</a> ]	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"> <li>• example: ONTAP_X</li> </ul>

## 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": "Switchover in progress.",
    "progress_summary": "MetroCluster updated successfully."
  },
  "nodes": {
    "name": "node1",
    "version": "ONTAP_X"
  },
  "pending_version": "ONTAP_X_1",
  "state": "completed",
  "status_details": {
    "action": "string",
    "end_time": "2019-02-02 19:00:00 UTC",
    "message": "Post-update checks successful",
    "name": "initialize",
    "node": {
      "name": "node1"
    },
    "start_time": "2019-02-02 19:00:00 UTC",
    "state": "failed"
  },
  "update_details": {
    "elapsed_duration": 2100,
    "estimated_duration": 4620,
    "node": {
      "name": "node1"
    },
    "phase": "Pre-update checks",
    "state": "failed"
  }
}
```

```
},
"validation_results": {
  "action": "string",
  "message": "string",
  "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	<a href="#">href</a>	

software\_mcc\_reference

Name	Type	Description
elapsed_duration	integer	Elapsed duration of update time (in seconds) in MetroCluster.
estimated_duration	integer	Estimated duration of update time (in seconds) in MetroCluster.
name	string	Name of the site in MetroCluster.
state		Upgrade state of MetroCluster.

metrocluster

Name	Type	Description
clusters	array[ <a href="#">software_mcc_reference</a> ]	List of MetroCluster sites, statuses, and active versions.
progress_details	string	MetroCluster update progress details.
progress_summary	string	MetroCluster update progress summary.

software\_node\_reference

Name	Type	Description
name	string	Name of the node.

Name	Type	Description
version	string	ONTAP version of the node. <ul style="list-style-type: none"> <li>• example: ONTAP_X</li> <li>• readOnly: 1</li> </ul>

node

Name	Type	Description
name	string	Name of the node to be retrieved for status details.

software\_status\_details\_reference

Name	Type	Description
action	string	Corrective action to be taken to resolve the status error.
end_time	string	End time for each status phase.
message	string	Detailed message of the phase details.
name	string	Name of the phase to be retrieved for status details.
node	<a href="#">node</a>	
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

Name	Type	Description
estimated_duration	integer	Estimated duration for each update phase
node	<a href="#">node</a>	
phase	string	Phase details
state	string	State of the update phase

#### software\_validation\_reference

Name	Type	Description
action	string	Corrective action to resolve errors or warnings for update checks.
message	string	Details of the error or warning encountered by the update check.
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[ <a href="#">error_arguments</a> ]	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

Upgrades the cluster software version. Setting `version` triggers the installation of the package to start. To validate the package for installation but not perform the installation, use the `validate_only` field on request. 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

### 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
<code>validate_only</code>	boolean	query	False	Validate the operation and its parameters, without actually performing the operation.
<code>skip_warnings</code>	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"> <li>enum: ["pause", "resume", "cancel"]</li> </ul>

#### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">metrocluster</a>	
nodes	array[ <a href="#">software_node_reference</a> ]	List of nodes and active versions.
pending_version	string	<p>Version being installed on the system.</p> <ul style="list-style-type: none"> <li>example: ONTAP_X_1</li> <li>readOnly: 1</li> </ul>
state	string	Operational state of the upgrade



Name	Type	Description
status_details	array[ <a href="#">software_status_details_reference</a> ]	Display status details.
update_details	array[ <a href="#">software_update_details_reference</a> ]	Display update process details.
validation_results	array[ <a href="#">software_validation_reference</a> ]	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"> <li>• example: ONTAP_X</li> </ul>

## Example request

```
{
  "_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": "Switchover in progress.",
    "progress_summary": "MetroCluster updated successfully."
  },
  "nodes": {
    "name": "node1",
    "version": "ONTAP_X"
  },
  "pending_version": "ONTAP_X_1",
  "state": "completed",
  "status_details": {
    "action": "string",
    "end_time": "2019-02-02 19:00:00 UTC",
    "message": "Post-update checks successful",
    "name": "initialize",
    "node": {
      "name": "node1"
    },
    "start_time": "2019-02-02 19:00:00 UTC",
    "state": "failed"
  },
  "update_details": {
    "elapsed_duration": 2100,
    "estimated_duration": 4620,
    "node": {
      "name": "node1"
    },
    "phase": "Pre-update checks",
    "state": "failed"
  }
}
```

```

    },
    "validation_results": {
      "action": "string",
      "message": "string",
      "status": "warning",
      "update_check": "nfs_mounts"
    },
    "version": "ONTAP_X"
  }
}

```

### 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	<a href="#">href</a>	

software\_mcc\_reference

Name	Type	Description
elapsed_duration	integer	Elapsed duration of update time (in seconds) in MetroCluster.
estimated_duration	integer	Estimated duration of update time (in seconds) in MetroCluster.
name	string	Name of the site in MetroCluster.
state		Upgrade state of MetroCluster.

metrocluster

Name	Type	Description
clusters	array[ <a href="#">software_mcc_reference</a> ]	List of MetroCluster sites, statuses, and active versions.
progress_details	string	MetroCluster update progress details.
progress_summary	string	MetroCluster update progress summary.

software\_node\_reference

Name	Type	Description
name	string	Name of the node.

Name	Type	Description
version	string	ONTAP version of the node. <ul style="list-style-type: none"> <li>• example: ONTAP_X</li> <li>• readOnly: 1</li> </ul>

node

Name	Type	Description
name	string	Name of the node to be retrieved for status details.

software\_status\_details\_reference

Name	Type	Description
action	string	Corrective action to be taken to resolve the status error.
end_time	string	End time for each status phase.
message	string	Detailed message of the phase details.
name	string	Name of the phase to be retrieved for status details.
node	<a href="#">node</a>	
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

Name	Type	Description
estimated_duration	integer	Estimated duration for each update phase
node	<a href="#">node</a>	
phase	string	Phase details
state	string	State of the update phase

#### software\_validation\_reference

Name	Type	Description
action	string	Corrective action to resolve errors or warnings for update checks.
message	string	Details of the error or warning encountered by the update check.
status	string	Status of this update check.
update_check	string	Name of the update check to be validated.

#### software\_reference

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">metrocluster</a>	
nodes	array[ <a href="#">software_node_reference</a> ]	List of nodes and active versions.

Name	Type	Description
pending_version	string	Version being installed on the system. <ul style="list-style-type: none"> <li>• example: ONTAP_X_1</li> <li>• readOnly: 1</li> </ul>
state	string	Operational state of the upgrade
status_details	array[ <a href="#">software_status_details_reference</a> ]	Display status details.
update_details	array[ <a href="#">software_update_details_reference</a> ]	Display update process details.
validation_results	array[ <a href="#">software_validation_reference</a> ]	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"> <li>• example: ONTAP_X</li> </ul>

#### job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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`

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_timeout	integer	query	False	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.

#### Request Body

Name	Type	Description
password	string	Password for download
url	string	HTTP or FTP URL of the package via 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": {
    "_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

### software\_package\_download

Name	Type	Description
password	string	Password for download
url	string	HTTP or FTP URL of the package via a server
username	string	Username for download

### href

Name	Type	Description
href	string	

### \_links

Name	Type	Description
self	<a href="#">href</a>	

### job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments

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

### Retrieve the software installation request history details

GET /cluster/software/history

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
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.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	
records	array[ <a href="#">software_history_reference</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "end_time": "2019-02-02 20:00:00 UTC",
    "from_version": "ONTAP_X1",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "start_time": "2019-02-02 19:00:00 UTC",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

software\_history\_reference

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"><li>• example: ONTAP_X1</li><li>• readOnly: 1</li></ul>
node	<a href="#">node</a>	
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"> <li>• example: ONTAP_X2</li> <li>• readOnly: 1</li> </ul>

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

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
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	
records	array[ <a href="#">software_package</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "create_time": "2019-02-04 19:00:00 UTC",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

software\_package

Name	Type	Description
_links	<a href="#">_links</a>	
create_time	string	Indicates when this package was loaded
version	string	Version of this package <ul style="list-style-type: none"><li>• example: ONTAP_X</li><li>• readOnly: 1</li></ul>

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code

Name	Type	Description
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}

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	

#### Response

Status: 202, Accepted

#### Error

Status: Default

#### 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 multiple nodes
10551325	Package does not exist on the system



Error codes	Description
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	<a href="#">error</a>	

### 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[ <a href="#">error_arguments</a> ]	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}

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	<a href="#">_links</a>	
create_time	string	Indicates when this package was loaded
version	string	Version of this package <ul style="list-style-type: none"><li>• example: ONTAP_X</li><li>• readOnly: 1</li></ul>

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "create_time": "2019-02-04 19:00:00 UTC",
  "version": "ONTAP_X"
}
```

### Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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.

## 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-logoff, 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

Retrieves audit configurations.

### Related ONTAP commands

- `vserver audit show`

### Learn more

- [DOC /protocols/audit](#)

### Parameters

Name	Type	In	Required	Description
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
events.security_group	boolean	query	False	Filter by events.security_group
events.file_share	boolean	query	False	Filter by events.file_share
events.file_operations	boolean	query	False	Filter by events.file_operations
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.cap_staging	boolean	query	False	Filter by events.cap_staging
log.format	string	query	False	Filter by log.format
log.rotation.schedule.hours	integer	query	False	Filter by log.rotation.schedule.hours
log.rotation.schedule.months	integer	query	False	Filter by log.rotation.schedule.months
log.rotation.schedule.minutes	integer	query	False	Filter by log.rotation.schedule.minutes
log.rotation.schedule.weekdays	integer	query	False	Filter by log.rotation.schedule.weekdays
log.rotation.schedule.days	integer	query	False	Filter by log.rotation.schedule.days



Name	Type	In	Required	Description
log.rotation.size	integer	query	False	Filter by log.rotation.size
log.retention.count	integer	query	False	Filter by log.retention.count
log.retention.duration	string	query	False	Filter by log.retention.duration
log_path	string	query	False	Filter by log_path
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">audit</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "log": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "format": "xml",
      "retention": {
        "duration": "P4DT12H30M5S"
      },
      "rotation": {
        "schedule": {
          "days": {
          },
          "hours": {
          },
          "minutes": {
          },
          "months": {
          },
          "weekdays": {
          }
        }
      }
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  }
}
```

```
}  
}
```

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

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	<a href="#">href</a>	

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. This is mutually exclusive with log size.

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	<a href="#">audit_schedule</a>	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. This is mutually exclusive with log size.
size	integer	Rotates logs based on log size in bytes. This is mutually exclusive with schedule.

#### log

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"> <li>• xml - Data ONTAP-specific XML log format</li> <li>• evtX - Microsoft Windows EVT X log format <ul style="list-style-type: none"> <li>◦ Default value: 1</li> <li>◦ enum: ["xml", "evtX"]</li> </ul> </li> </ul>
retention	<a href="#">retention</a>	
rotation	<a href="#">rotation</a>	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

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">events</a>	
log	<a href="#">log</a>	
log_path	string	The audit log destination path where consolidated audit logs are stored.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message



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

## Create an audit configuration

POST /protocols/audit

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](#)

## Request Body

Name	Type	Description
enabled	boolean	Specifies whether or not auditing is enabled on the SVM.
events	<a href="#">events</a>	
log	<a href="#">log</a>	
log_path	string	The audit log destination path where consolidated audit logs are stored.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "log": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "format": "xml",
    "retention": {
      "duration": "P4DT12H30M5S"
    },
    "rotation": {
      "schedule": {
        "days": {
        },
        "hours": {
        },
        "minutes": {
        },
        "months": {
        },
        "weekdays": {
        }
      }
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

## Response

Status: 202, Accepted

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">audit</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "log": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "format": "xml",
      "retention": {
        "duration": "P4DT12H30M5S"
      },
      "rotation": {
        "schedule": {
          "days": {
          },
          "hours": {
          },
          "minutes": {
          },
          "months": {
          },
          "weekdays": {
          }
        }
      }
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "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

Error Code	Description
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
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

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	<a href="#">href</a>	

### 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. This is mutually exclusive with log size.

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	<a href="#">audit_schedule</a>	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. This is mutually exclusive with log size.
size	integer	Rotates logs based on log size in bytes. This is mutually exclusive with schedule.

## log

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"> <li>• xml - Data ONTAP-specific XML log format</li> <li>• evtX - Microsoft Windows EVT X log format <ul style="list-style-type: none"> <li>◦ Default value: 1</li> <li>◦ enum: ["xml", "evtX"]</li> </ul> </li> </ul>
retention	<a href="#">retention</a>	
rotation	<a href="#">rotation</a>	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

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">events</a>	
log	<a href="#">log</a>	
log_path	string	The audit log destination path where consolidated audit logs are stored.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments

Name	Type	Description
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}

Deletes an audit configuration.

#### Related ONTAP commands

- `vserver audit disable`
- `vserver audit delete`

#### 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.

#### Response

Status: 202, Accepted

#### Error

Status: Default

#### ONTAP Error Response Codes

Error Code	Description
9699349	Auditing should be disabled before deleting the audit configuration

Error Code	Description
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve the audit configuration for an SVM

GET /protocols/audit/{svm.uuid}

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	<a href="#">events</a>	
log	<a href="#">log</a>	
log_path	string	The audit log destination path where consolidated audit logs are stored.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example response

```
{
  "log": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "format": "xml",
    "retention": {
      "duration": "P4DT12H30M5S"
    },
    "rotation": {
      "schedule": {
        "days": {
        },
        "hours": {
        },
        "minutes": {
        },
        "months": {
        },
        "weekdays": {
        }
      }
    }
  },
  "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	<a href="#">href</a>	

### 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. This is mutually exclusive with log size.

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	<a href="#">audit_schedule</a>	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. This is mutually exclusive with log size.
size	integer	Rotates logs based on log size in bytes. This is mutually exclusive with schedule.

## log

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"> <li>• xml - Data ONTAP-specific XML log format</li> <li>• evtx - Microsoft Windows EVTX log format <ul style="list-style-type: none"> <li>◦ Default value: 1</li> <li>◦ enum: ["xml", "evtx"]</li> </ul> </li> </ul>
retention	<a href="#">retention</a>	
rotation	<a href="#">rotation</a>	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

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the audit configuration for an SVM

PATCH /protocols/audit/{svm.uuid}

Updates an audit configuration for an SVM.

### Related ONTAP commands

- `vserver audit modify`

### 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.

#### Request Body

Name	Type	Description
enabled	boolean	Specifies whether or not auditing is enabled on the SVM.
events	<a href="#">events</a>	
log	<a href="#">log</a>	
log_path	string	The audit log destination path where consolidated audit logs are stored.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "log": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "format": "xml",
    "retention": {
      "duration": "P4DT12H30M5S"
    },
    "rotation": {
      "schedule": {
        "days": {
        },
        "hours": {
        },
        "minutes": {
        },
        "months": {
        },
        "weekdays": {
        }
      }
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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
9699409	Failed to enable multiproto.audit.evtxlog.support support capability
9699410	Failed to disable multiproto.audit.evtxlog.support support capability



Error Code	Description
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

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	<a href="#">href</a>	

### 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. This is mutually exclusive with log size.

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	<a href="#">audit_schedule</a>	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. This is mutually exclusive with log size.
size	integer	Rotates logs based on log size in bytes. This is mutually exclusive with schedule.

## log

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"> <li>• xml - Data ONTAP-specific XML log format</li> <li>• evtx - Microsoft Windows EVTX log format <ul style="list-style-type: none"> <li>◦ Default value: 1</li> <li>◦ enum: ["xml", "evtx"]</li> </ul> </li> </ul>
retention	<a href="#">retention</a>	
rotation	<a href="#">rotation</a>	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

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">events</a>	
log	<a href="#">log</a>	
log_path	string	The audit log destination path where consolidated audit logs are stored.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

## error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

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.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">cifs_search_path</a> ]	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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.

Name	Type	Description
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

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](#)

### 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	SVM, applies only to SVM-scoped objects.

### Example request

```
{
  "index": 0,
  "path": "/HomeDirectory/EngDomain",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

### Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">cifs_search_path</a> ]	

## 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

## ONTAP Error Response Codes

Error Code	Description
655551	Invalid home-directory search-path path
655462	The specified path is an invalid file-type

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

\_links

Name	Type	Description
next	<a href="#">href</a>	

Name	Type	Description
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
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}
```

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
index	integer	path	True	Home directory search path index

## 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[ <a href="#">error_arguments</a> ]	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}
```

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
index	integer	path	True	Home directory search path index

Name	Type	In	Required	Description
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	SVM, applies only to SVM-scoped objects.

### 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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
index	integer	path	True	Home directory search path index
new_index	integer	query	False	New position for the home directory search path

### 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

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	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.

## 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 `<em>multichannel</em>` 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, \"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,  
        "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,
  "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"
}
],
"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,
    "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"
}
```

## 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
```

Retrieves CIFS servers.

### 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
comment	string	query	False	Filter by comment
security.smb_encryption	boolean	query	False	Filter by security.smb_encryption
security.smb_signing	boolean	query	False	Filter by security.smb_signing
security.restrict_anonymous	string	query	False	Filter by security.restrict_anonymous
security.kdc_encryption	boolean	query	False	Filter by security.kdc_encryption
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
name	string	query	False	Filter by name
default_unix_user	string	query	False	Filter by default_unix_user
netbios.enabled	boolean	query	False	Filter by netbios.enabled
netbios.aliases	string	query	False	Filter by netbios.aliases
netbios.wins_servers	string	query	False	Filter by netbios.wins_servers
ad_domain.organizational_unit	string	query	False	Filter by ad_domain.organizational_unit
ad_domain.fqdn	string	query	False	Filter by ad_domain.fqdn

Name	Type	In	Required	Description
ad_domain.user	string	query	False	Filter by ad_domain.user
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	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">cifs_service</a> ]	

**Example response**

A large, empty rectangular box with a thin, dashed border, occupying most of the page. It is intended for an example response.

```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ad_domain": {
      "fqdn": "example.com"
    },
    "comment": "This CIFS Server Belongs to CS Department",
    "name": "CIFS1",
    "netbios": {
      "aliases": [
        "ALIAS_1",
        "ALIAS_2",
        "ALIAS_3"
      ],
      "wins_servers": [
        "10.224.65.20",
        "10.224.65.21"
      ]
    },
    "security": {
      "restrict_anonymous": "no_restriction"
    },
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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\_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
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. 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"> <li>• RC4-HMAC</li> <li>• DES</li> <li>• AES</li> </ul> <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	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: <ul style="list-style-type: none"> <li>no_restriction - No access restriction for an anonymous user.</li> <li>no_enumeration - Enumeration is restricted for an anonymous user.</li> <li>no_access - All access is restricted for an anonymous user.</li> </ul>
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.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cifs\_service

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>ad_domain</code>	<a href="#">ad_domain</a>	
<code>comment</code>	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.
<code>default_unix_user</code>	string	Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails.
<code>enabled</code>	boolean	Specifies if the CIFS service is administratively enabled.
<code>name</code>	string	The name of the CIFS server.
<code>netbios</code>	<a href="#">cifs_netbios</a>	
<code>security</code>	<a href="#">cifs_service_security</a>	
<code>svm</code>	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### 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[ <a href="#">error_arguments</a> ]	Message arguments
<code>code</code>	string	Error code
<code>message</code>	string	Error message

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

### Create a CIFS server

POST /protocols/cifs/services

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*
- `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](#)

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>ad_domain</code>	<a href="#">ad_domain</a>	
<code>comment</code>	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.
<code>default_unix_user</code>	string	Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails.
<code>enabled</code>	boolean	Specifies if the CIFS service is administratively enabled.
<code>name</code>	string	The name of the CIFS server.
<code>netbios</code>	<a href="#">cifs_netbios</a>	
<code>security</code>	<a href="#">cifs_service_security</a>	
<code>svm</code>	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_domain": {
    "fqdn": "example.com"
  },
  "comment": "This CIFS Server Belongs to CS Department",
  "name": "CIFS1",
  "netbios": {
    "aliases": [
      "ALIAS_1",
      "ALIAS_2",
      "ALIAS_3"
    ],
    "wins_servers": [
      "10.224.65.20",
      "10.224.65.21"
    ]
  },
  "security": {
    "restrict_anonymous": "no_restriction"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

## 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	<a href="#">href</a>	

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\_netbios

Name	Type	Description
aliases	array[string]	

Name	Type	Description
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
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. 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"> <li>• RC4-HMAC</li> <li>• DES</li> <li>• AES</li> </ul> <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	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: <ul style="list-style-type: none"> <li>no_restriction - No access restriction for an anonymous user.</li> <li>no_enumeration - Enumeration is restricted for an anonymous user.</li> <li>no_access - All access is restricted for an anonymous user.</li> </ul>
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.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cifs\_service

Name	Type	Description
_links	<a href="#">_links</a>	
ad_domain	<a href="#">ad_domain</a>	
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.
name	string	The name of the CIFS server.
netbios	<a href="#">cifs_netbios</a>	
security	<a href="#">cifs_service_security</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}
```

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

#### Request Body

Name	Type	Description
ad_domain	<a href="#">ad_domain</a>	

### Example request

```
{
  "ad_domain": {
    "fqdn": "example.com"
  }
}
```

### 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	<a href="#">ad_domain</a>	

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	

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[ <a href="#">error_arguments</a> ]	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}

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	



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

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
<a href="#">ad_domain</a>	<a href="#">ad_domain</a>	
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.
name	string	The name of the CIFS server.
netbios	<a href="#">cifs_netbios</a>	
security	<a href="#">cifs_service_security</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_domain": {
    "fqdn": "example.com"
  },
  "comment": "This CIFS Server Belongs to CS Department",
  "name": "CIFS1",
  "netbios": {
    "aliases": [
      "ALIAS_1",
      "ALIAS_2",
      "ALIAS_3"
    ],
    "wins_servers": [
      "10.224.65.20",
      "10.224.65.21"
    ]
  },
  "security": {
    "restrict_anonymous": "no_restriction"
  },
  "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	<a href="#">href</a>	

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\_netbios

Name	Type	Description
aliases	array[string]	

Name	Type	Description
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
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. 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"> <li>• RC4-HMAC</li> <li>• DES</li> <li>• AES</li> </ul> <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	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: <ul style="list-style-type: none"> <li>no_restriction - No access restriction for an anonymous user.</li> <li>no_enumeration - Enumeration is restricted for an anonymous user.</li> <li>no_access - All access is restricted for an anonymous user.</li> </ul>
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.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.



**Request Body**

<b>Name</b>	<b>Type</b>	<b>Description</b>
_links	<a href="#">_links</a>	
ad_domain	<a href="#">ad_domain</a>	
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.
name	string	The name of the CIFS server.
netbios	<a href="#">cifs_netbios</a>	
security	<a href="#">cifs_service_security</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_domain": {
    "fqdn": "example.com"
  },
  "comment": "This CIFS Server Belongs to CS Department",
  "name": "CIFS1",
  "netbios": {
    "aliases": [
      "ALIAS_1",
      "ALIAS_2",
      "ALIAS_3"
    ],
    "wins_servers": [
      "10.224.65.20",
      "10.224.65.21"
    ]
  },
  "security": {
    "restrict_anonymous": "no_restriction"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

## 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	<a href="#">href</a>	

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\_netbios

Name	Type	Description
aliases	array[string]	

Name	Type	Description
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
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. 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"> <li>• RC4-HMAC</li> <li>• DES</li> <li>• AES</li> </ul> <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	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: <ul style="list-style-type: none"> <li>no_restriction - No access restriction for an anonymous user.</li> <li>no_enumeration - Enumeration is restricted for an anonymous user.</li> <li>no_access - All access is restricted for an anonymous user.</li> </ul>
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.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cifs\_service

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
<a href="#">ad_domain</a>	<a href="#">ad_domain</a>	
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.
name	string	The name of the CIFS server.
netbios	<a href="#">cifs_netbios</a>	
security	<a href="#">cifs_service_security</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### job\_link

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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
```

Retrieves CIFS shares.

### Related ONTAP commands

- `vserver cifs share show`
- `vserver cifs share properties show`

### Learn more

- [DOC /protocols/cifs/shares](#)

### Parameters

<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
encryption	boolean	query	False	Filter by encryption
change_notify	boolean	query	False	Filter by change_notify
path	string	query	False	Filter by path
comment	string	query	False	Filter by comment
unix_symlink	string	query	False	Filter by unix_symlink
oplocks	boolean	query	False	Filter by oplocks
access_based_enumeration	boolean	query	False	Filter by access_based_enumeration
home_directory	boolean	query	False	Filter by home_directory
name	string	query	False	Filter by name
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
volume.name	string	query	False	Filter by volume.name
volume.uuid	string	query	False	Filter by volume.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.

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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">cifs_share</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "acls": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "permission": "no_access",
      "type": "windows",
      "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": "local",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

cifs\_share\_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• no_access - User does not have CIFS share access</li><li>• read - User has only read access</li><li>• change - User has change access</li><li>• full_control - User has full_control access</li></ul>

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"> <li>• windows - Windows user or group</li> <li>• unix_user - UNIX user</li> <li>• unix_group - UNIX group</li> </ul>
user_or_group	string	Specifies the user or group name to add to the access control list of a CIFS share.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

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
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">cifs_share_acl</a> ]	
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"> <li>• Default value:</li> <li>• readCreate: 1</li> </ul>
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"> <li>• example: /volume_1/eng_vol/</li> <li>• maxLength: 256</li> <li>• minLength: 1</li> </ul>
svm	svm	SVM, applies only to SVM-scoped objects.
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"> <li>• local - Enables only local symbolic links which is within the same CIFS share.</li> <li>• widelink - Enables both local symlinks and widelinks.</li> <li>• disable - Disables local symlinks and widelinks.</li> </ul>
volume	volume	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

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](#)

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>access_based_enumeration</code>	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.
<code>acls</code>	array[ <a href="#">cifs_share_acl</a> ]	
<code>change_notify</code>	boolean	Specifies whether CIFS clients can request for change notifications for directories on this share.
<code>comment</code>	string	Specify the CIFS share descriptions.
<code>encryption</code>	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"> <li>• Default value: 1</li> <li>• readCreate: 1</li> </ul>
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. 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"> <li>• example: /volume_1/eng_vol/</li> <li>• maxLength: 256</li> <li>• minLength: 1</li> </ul>
svm	svm	SVM, applies only to SVM-scoped objects.
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"> <li>• local - Enables only local symbolic links which is within the same CIFS share.</li> <li>• widelink - Enables both local symlinks and widelinks.</li> <li>• disable - Disables local symlinks and widelinks.</li> </ul>
volume	volume	

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "acls": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "permission": "no_access",
    "type": "windows",
    "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": "local",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volumel",
    "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
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

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	<a href="#">href</a>	

cifs\_share\_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• no_access - User does not have CIFS share access</li><li>• read - User has only read access</li><li>• change - User has change access</li><li>• full_control - User has full_control access</li></ul>
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"><li>• windows - Windows user or group</li><li>• unix_user - UNIX user</li><li>• unix_group - UNIX group</li></ul>
user_or_group	string	Specifies the user or group name to add to the access control list of a CIFS share.

## svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

## volume

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"><li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li></ul>

## 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
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">cifs_share_acl</a> ]	

Name	Type	Description
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.
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"> <li>• Default value: 1</li> <li>• readCreate: 1</li> </ul>

Name	Type	Description
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>
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"> <li>• example: /volume_1/eng_vol/</li> <li>• maxLength: 256</li> <li>• minLength: 1</li> </ul>
svm	svm	<p>SVM, applies only to SVM-scoped objects.</p>

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"> <li>• local - Enables only local symbolic links which is within the same CIFS share.</li> <li>• widelink - Enables both local symlinks and widelinks.</li> <li>• disable - Disables local symlinks and widelinks.</li> </ul>
volume	<a href="#">volume</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}
```

Deletes a CIFS share.

#### Related ONTAP commands

- `vserver cifs share delete`

## Learn more

- [DOC /protocols/cifs/shares](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	Share Name

## 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	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	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}

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.



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

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">cifs_share_acl</a> ]	
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"> <li>• Default value: 1</li> <li>• readCreate: 1</li> </ul>
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. 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"> <li>• example: /volume_1/eng_vol/</li> <li>• maxLength: 256</li> <li>• minLength: 1</li> </ul>
svm	svm	SVM, applies only to SVM-scoped objects.
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"> <li>• local - Enables only local symbolic links which is within the same CIFS share.</li> <li>• widelink - Enables both local symlinks and widelinks.</li> <li>• disable - Disables local symlinks and widelinks.</li> </ul>
volume	volume	

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "acls": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "permission": "no_access",
    "type": "windows",
    "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": "local",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volumel",
    "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	<a href="#">href</a>	

cifs\_share\_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• no_access - User does not have CIFS share access</li><li>• read - User has only read access</li><li>• change - User has change access</li><li>• full_control - User has full_control access</li></ul>
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"><li>• windows - Windows user or group</li><li>• unix_user - UNIX user</li><li>• unix_group - UNIX group</li></ul>
user_or_group	string	Specifies the user or group name to add to the access control list of a CIFS share.

## svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

## volume

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"><li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li></ul>

## error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

## error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	Share Name

### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">cifs_share_acl</a> ]	
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"> <li>• Default value: 1</li> <li>• readCreate: 1</li> </ul>
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.

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"> <li>• example: /volume_1/eng_vol/</li> <li>• maxLength: 256</li> <li>• minLength: 1</li> </ul>
svm	svm	SVM, applies only to SVM-scoped objects.
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"> <li>• local - Enables only local symbolic links which is within the same CIFS share.</li> <li>• widelink - Enables both local symlinks and widelinks.</li> <li>• disable - Disables local symlinks and widelinks.</li> </ul>
volume	volume	

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "acls": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "permission": "no_access",
    "type": "windows",
    "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": "local",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volumel",
    "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
656425	Failed to modify the CIFS share because the path for an administrative share cannot be modified

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

cifs\_share\_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• no_access - User does not have CIFS share access</li><li>• read - User has only read access</li><li>• change - User has change access</li><li>• full_control - User has full_control access</li></ul>
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"><li>• windows - Windows user or group</li><li>• unix_user - UNIX user</li><li>• unix_group - UNIX group</li></ul>
user_or_group	string	Specifies the user or group name to add to the access control list of a CIFS share.

## svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

## volume

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"><li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li></ul>

## 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
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">cifs_share_acl</a> ]	

Name	Type	Description
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.
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"> <li>• Default value: 1</li> <li>• readCreate: 1</li> </ul>



Name	Type	Description
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>
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"> <li>• example: /volume_1/eng_vol/</li> <li>• maxLength: 256</li> <li>• minLength: 1</li> </ul>
svm	svm	<p>SVM, applies only to SVM-scoped objects.</p>

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"> <li>• local - Enables only local symbolic links which is within the same CIFS share.</li> <li>• widelink - Enables both local symlinks and widelinks.</li> <li>• disable - Disables local symlinks and widelinks.</li> </ul>
volume	<a href="#">volume</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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 acs 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

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
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
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">cifs_share_acl</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "permission": "no_access",
    "type": "windows",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

cifs\_share\_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• no_access - User does not have CIFS share access</li><li>• read - User has only read access</li><li>• change - User has change access</li><li>• full_control - User has full_control access</li></ul>

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"> <li>• windows - Windows user or group</li> <li>• unix_user - UNIX user</li> <li>• unix_group - UNIX group</li> </ul>
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[ <a href="#">error_arguments</a> ]	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

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
share	string	path	True	CIFS Share Name

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>permission</code>	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"> <li>• <code>no_access</code> - User does not have CIFS share access</li> <li>• <code>read</code> - User has only read access</li> <li>• <code>change</code> - User has change access</li> <li>• <code>full_control</code> - User has full_control access</li> </ul>

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"> <li>• windows - Windows user or group</li> <li>• unix_user - UNIX user</li> <li>• unix_group - UNIX group</li> </ul>
user_or_group	string	Specifies the user or group name to add to the access control list of a CIFS share.

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "permission": "no_access",
  "type": "windows",
  "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

Error Code	Description
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

Name	Type	Description
error	<a href="#">error</a>	

### 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	<a href="#">href</a>	

cifs\_share\_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• no_access - User does not have CIFS share access</li><li>• read - User has only read access</li><li>• change - User has change access</li><li>• full_control - User has full_control access</li></ul>
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"><li>• windows - Windows user or group</li><li>• unix_user - UNIX user</li><li>• unix_group - UNIX group</li></ul>
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[ <a href="#">error_arguments</a> ]	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}
```

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
share	string	path	True	Share name
user_or_group	string	path	True	User or group name
type	string	path	True	User or group type



## 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[ <a href="#">error_arguments</a> ]	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}
```

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
share	string	path	True	Share name
user_or_group	string	path	True	User or group name

Name	Type	In	Required	Description
type	string	path	True	User or group type
fields	array[string]	query	False	Specify the fields to return.

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
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"> <li>• no_access - User does not have CIFS share access</li> <li>• read - User has only read access</li> <li>• change - User has change access</li> <li>• full_control - User has full_control access</li> </ul>
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"> <li>• windows - Windows user or group</li> <li>• unix_user - UNIX user</li> <li>• unix_group - UNIX group</li> </ul>
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": "no_access",
  "type": "windows",
  "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	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}
```

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
share	string	path	True	Share name
user_or_group	string	path	True	User or group name
type	string	path	True	User or group type

## Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"><li>• no_access - User does not have CIFS share access</li><li>• read - User has only read access</li><li>• change - User has change access</li><li>• full_control - User has full_control access</li></ul>
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"><li>• windows - Windows user or group</li><li>• unix_user - UNIX user</li><li>• unix_group - UNIX group</li></ul>
user_or_group	string	Specifies the user or group name to add to the access control list of a CIFS share.

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "permission": "no_access",
  "type": "windows",
  "user_or_group": "ENGDOMAIN\\ad_user"
}
```

### 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

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	<a href="#">href</a>	

cifs\_share\_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• no_access - User does not have CIFS share access</li><li>• read - User has only read access</li><li>• change - User has change access</li><li>• full_control - User has full_control access</li></ul>
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"><li>• windows - Windows user or group</li><li>• unix_user - UNIX user</li><li>• unix_group - UNIX group</li></ul>
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
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

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
target.share	string	query	False	Filter by target.share
target.locality	string	query	False	Filter by target.locality
target.server	string	query	False	Filter by target.server
target.path	string	query	False	Filter by target.path
target.home_directory	boolean	query	False	Filter by target.home_directory
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
unix_path	string	query	False	Filter by unix_path

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	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">cifs_symlink_mapping</a> ]	



## 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": "local",
      "path": "/dir1/dir2/",
      "server": "ENG_CIFS",
      "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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.

Name	Type	Description
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"> <li>• local - Local symbolic link maps only to the same CIFS share.</li> <li>• widelink - Wide symbolic link maps to any CIFS share on the network.</li> </ul>
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.</p> <p>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>
server	string	<p>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'.</p> <p>You can specify the value in any of the following formats:</p> <ul style="list-style-type: none"> <li>• DNS name of the CIFS server.</li> <li>• IP address of the CIFS server.</li> <li>• NetBIOS name of the CIFS server.</li> </ul>
share	string	<p>Specifies the CIFS share name on the destination CIFS server to which the UNIX symbolic link is pointing.</p>

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	<a href="#">_links</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">cifs_target</a>	
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[ <a href="#">error_arguments</a> ]	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

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](#)

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>svm</code>	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
<code>target</code>	<a href="#">cifs_target</a>	
<code>unix_path</code>	string	Specifies the UNIX path prefix to be matched for the mapping.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "locality": "local",
    "path": "/dir1/dir2/",
    "server": "ENGCIFS",
    "share": "ENG_SHARE"
  },
  "unix_path": "/mnt/eng_volume/"
}
```

## Response

Status: 201, Created

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records
records	<a href="#">array[cifs_symlink_mapping]</a>	

## 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": "local",
      "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
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

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	Specifies whether the CIFS symbolic link is a local link or wide link. The following values are supported: <ul style="list-style-type: none"><li>• local - Local symbolic link maps only to the same CIFS share.</li><li>• widelink - Wide symbolic link maps to any CIFS share on the network.</li></ul>

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"> <li>• DNS name of the CIFS server.</li> <li>• IP address of the CIFS server.</li> <li>• NetBIOS name of the CIFS server.</li> </ul>
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	<a href="#">_links</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">cifs_target</a>	

Name	Type	Description
unix_path	string	Specifies the UNIX path prefix to be matched for the mapping.

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete a UNIX symbolic link mapping for CIFS clients

```
DELETE /protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}
```

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
unix_path	string	path	True	UNIX symbolic link path

## Response

Status: 200, Ok

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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 UNIX symbolic link mapping for CIFS clients

GET /protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix\_path}

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
unix_path	string	path	True	UNIX symbolic link path

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

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
<a href="#">svm</a>	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
<a href="#">target</a>	<a href="#">cifs_target</a>	
<a href="#">unix_path</a>	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": "local",
    "path": "/dir1/dir2/",
    "server": "ENGCIFS",
    "share": "ENG_SHARE"
  },
  "unix_path": "/mnt/eng_volume/"
}
```

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	



## 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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	Specifies whether the CIFS symbolic link is a local link or wide link. The following values are supported: <ul style="list-style-type: none"><li>• local - Local symbolic link maps only to the same CIFS share.</li><li>• widelink - Wide symbolic link maps to any CIFS share on the network.</li></ul>

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"> <li>• DNS name of the CIFS server.</li> <li>• IP address of the CIFS server.</li> <li>• NetBIOS name of the CIFS server.</li> </ul>
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[ <a href="#">error_arguments</a> ]	Message arguments

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

### Update a UNIX symbolic link mapping for CIFS clients

PATCH /protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix\_path}

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
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
unix_path	string	path	True	UNIX symbolic link path

#### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">cifs_target</a>	
unix_path	string	Specifies the UNIX path prefix to be matched for the mapping.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "locality": "local",
    "path": "/dir1/dir2/",
    "server": "ENG_CIFS",
    "share": "ENG_SHARE"
  },
  "unix_path": "/mnt/eng_volume/"
}
```

## 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

Error Code	Description
655547	Failed to modify symlink mapping because administrative share cannot be used as target share

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	Specifies whether the CIFS symbolic link is a local link or wide link. The following values are supported: <ul style="list-style-type: none"><li>• local - Local symbolic link maps only to the same CIFS share.</li><li>• widelink - Wide symbolic link maps to any CIFS share on the network.</li></ul>

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"> <li>• DNS name of the CIFS server.</li> <li>• IP address of the CIFS server.</li> <li>• NetBIOS name of the CIFS server.</li> </ul>
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	<a href="#">_links</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">cifs_target</a>	



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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage FPolicy configuration

### 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\": [
\"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": "pol10",
  "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"

```

## Retrieve an FPolicy configuration

```
GET /protocols/fpolicy
```

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

- [DOC /protocols/fpolicy](#)

#### Parameters

Name	Type	In	Required	Description
<code>engines.primary_servers</code>	string	query	False	Filter by <code>engines.primary_servers</code>
<code>engines.port</code>	integer	query	False	Filter by <code>engines.port</code>
<code>engines.type</code>	string	query	False	Filter by <code>engines.type</code>
<code>engines.secondary_servers</code>	string	query	False	Filter by <code>engines.secondary_servers</code>
<code>engines.name</code>	string	query	False	Filter by <code>engines.name</code>
<code>events.name</code>	string	query	False	Filter by <code>events.name</code>
<code>events.filters.first_write</code>	boolean	query	False	Filter by <code>events.filters.first_write</code>
<code>events.filters.setattr_with_size_change</code>	boolean	query	False	Filter by <code>events.filters.setattr_with_size_change</code>
<code>events.filters.monitor_ads</code>	boolean	query	False	Filter by <code>events.filters.monitor_ads</code>
<code>events.filters.close_with_read</code>	boolean	query	False	Filter by <code>events.filters.close_with_read</code>

<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
events.filters.setattr_with_group_change	boolean	query	False	Filter by events.filters.setattr_with_group_change
events.filters.offline_bit	boolean	query	False	Filter by events.filters.offline_bit
events.filters.setattr_with_sacl_change	boolean	query	False	Filter by events.filters.setattr_with_sacl_change
events.filters.setattr_with_dacl_change	boolean	query	False	Filter by events.filters.setattr_with_dacl_change
events.filters.open_with_write_intent	boolean	query	False	Filter by events.filters.open_with_write_intent
events.filters.setattr_with_modify_time_change	boolean	query	False	Filter by events.filters.setattr_with_modify_time_change
events.filters.setattr_with_creation_time_change	boolean	query	False	Filter by events.filters.setattr_with_creation_time_change
events.filters.setattr_with_access_time_change	boolean	query	False	Filter by events.filters.setattr_with_access_time_change
events.filters.open_with_delete_intent	boolean	query	False	Filter by events.filters.open_with_delete_intent
events.filters.setattr_with_allocation_size_change	boolean	query	False	Filter by events.filters.setattr_with_allocation_size_change
events.filters.close_without_modification	boolean	query	False	Filter by events.filters.close_without_modification

Name	Type	In	Required	Description
events.filters.write_with_size_change	boolean	query	False	Filter by events.filters.write_with_size_change
events.filters.close_with_modification	boolean	query	False	Filter by events.filters.close_with_modification
events.filters.exclude_directory	boolean	query	False	Filter by events.filters.exclude_directory
events.filters.setattr_with_mode_change	boolean	query	False	Filter by events.filters.setattr_with_mode_change
events.filters.first_read	boolean	query	False	Filter by events.filters.first_read
events.filters.setattr_with_owner_change	boolean	query	False	Filter by events.filters.setattr_with_owner_change
events.protocol	string	query	False	Filter by events.protocol
events.volume_monitoring	boolean	query	False	Filter by events.volume_monitoring
events.file_operations.link	boolean	query	False	Filter by events.file_operations.link
events.file_operations.create	boolean	query	False	Filter by events.file_operations.create
events.file_operations.close	boolean	query	False	Filter by events.file_operations.close
events.file_operations.setattr	boolean	query	False	Filter by events.file_operations.setattr

<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
events.file_operation s.rename	boolean	query	False	Filter by events.file_operation s.rename
events.file_operation s.delete	boolean	query	False	Filter by events.file_operation s.delete
events.file_operation s.read	boolean	query	False	Filter by events.file_operation s.read
events.file_operation s.lookup	boolean	query	False	Filter by events.file_operation s.lookup
events.file_operation s.getattr	boolean	query	False	Filter by events.file_operation s.getattr
events.file_operation s.create_dir	boolean	query	False	Filter by events.file_operation s.create_dir
events.file_operation s.rename_dir	boolean	query	False	Filter by events.file_operation s.rename_dir
events.file_operation s.open	boolean	query	False	Filter by events.file_operation s.open
events.file_operation s.delete_dir	boolean	query	False	Filter by events.file_operation s.delete_dir
events.file_operation s.write	boolean	query	False	Filter by events.file_operation s.write
events.file_operation s.symlink	boolean	query	False	Filter by events.file_operation s.symlink
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name

Name	Type	In	Required	Description
policies.events.name	string	query	False	Filter by policies.events.name
policies.priority	integer	query	False	Filter by policies.priority
policies.mandatory	boolean	query	False	Filter by policies.mandatory
policies.engine.name	string	query	False	Filter by policies.engine.name
policies.scope.include_shares	string	query	False	Filter by policies.scope.include_shares
policies.scope.include_export_policies	string	query	False	Filter by policies.scope.include_export_policies
policies.scope.include_volumes	string	query	False	Filter by policies.scope.include_volumes
policies.scope.exclude_export_policies	string	query	False	Filter by policies.scope.exclude_export_policies
policies.scope.include_extension	string	query	False	Filter by policies.scope.include_extension
policies.scope.exclude_shares	string	query	False	Filter by policies.scope.exclude_shares
policies.scope.exclude_extension	string	query	False	Filter by policies.scope.exclude_extension
policies.scope.exclude_volumes	string	query	False	Filter by policies.scope.exclude_volumes

Name	Type	In	Required	Description
policies.name	string	query	False	Filter by policies.name
policies.enabled	boolean	query	False	Filter by policies.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	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">fpolicy</a> ]	

## Example response

```
{
  "_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": "synchronous"
    },
    "events": {
      "name": "event_nfs_close",
      "protocol": "cifs"
    },
    "policies": {
      "engine": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      }
    },
    "events": [
      "event_nfs_close",
      "event_open"
    ],
  },
}
```

```
"name": "fp_policy_1",
"scope": {
  "exclude_export_policies": {
  },
  "exclude_extension": {
  },
  "exclude_shares": {
  },
  "exclude_volumes": [
    "voll",
    "vol_svm1",
    "*"
  ],
  "include_export_policies": {
  },
  "include_extension": {
  },
  "include_shares": [
    "sh1",
    "share_cifs"
  ],
  "include_volumes": [
    "voll",
    "vol_svm1"
  ]
},
"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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

fpolicy\_engine

The engine defines how ONTAP makes and manages connections to external FPolicy servers.

Name	Type	Description
name	string	Specifies the name to assign to the external server configuration.
port	integer	Port number of the FPolicy server application.
primary_servers	array[string]	
secondary_servers	array[string]	

Name	Type	Description
type	string	<p>The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:</p> <ul style="list-style-type: none"> <li>• synchronous - After sending a notification, wait for a response from the FPolicy server.</li> <li>• asynchronous - After sending a notification, file request processing continues. <ul style="list-style-type: none"> <li>◦ Default value: 1</li> <li>◦ enum: ["synchronous", "asynchronous"]</li> </ul> </li> </ul>

### 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.

Name	Type	Description
close	boolean	File close operations
create	boolean	File create operations
create_dir	boolean	Directory create operations
delete	boolean	File delete operations
delete_dir	boolean	Directory delete operations
getattr	boolean	Get attribute operations
link	boolean	Link operations
lookup	boolean	Lookup operations
open	boolean	File open operations
read	boolean	File read operations
rename	boolean	File rename operations

Name	Type	Description
rename_dir	boolean	Directory rename operations
setattr	boolean	Set attribute operations
symlink	boolean	Symbolic link operations
write	boolean	File write operations

## 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.

Name	Type	Description
close_with_modification	boolean	Filter the client request for close with modification.
close_with_read	boolean	Filter the client request for close with read.
close_without_modification	boolean	Filter the client request for close without modification.
exclude_directory	boolean	Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.
first_read	boolean	Filter the client requests for the first-read.
first_write	boolean	Filter the client requests for the first-write.
monitor_ads	boolean	Filter the client request for alternate data stream.
offline_bit	boolean	Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.
open_with_delete_intent	boolean	Filter the client request for open with delete intent.

Name	Type	Description
open_with_write_intent	boolean	Filter the client request for open with write intent.
setattr_with_access_time_change	boolean	Filter the client setattr requests for changing the access time of a file or directory.
setattr_with_allocation_size_change	boolean	Filter the client setattr requests for changing the allocation size of a file.
setattr_with_creation_time_change	boolean	Filter the client setattr requests for changing the creation time of a file or directory.
setattr_with_dacl_change	boolean	Filter the client setattr requests for changing dacl on a file or directory.
setattr_with_group_change	boolean	Filter the client setattr requests for changing group of a file or directory.
setattr_with_mode_change	boolean	Filter the client setattr requests for changing the mode bits on a file or directory.
setattr_with_modify_time_change	boolean	Filter the client setattr requests for changing the modification time of a file or directory.
setattr_with_owner_change	boolean	Filter the client setattr requests for changing owner of a file or directory.
setattr_with_sacl_change	boolean	Filter the client setattr requests for changing sacl on a file or directory.
setattr_with_size_change	boolean	Filter the client setattr requests for changing the size of a file.
write_with_size_change	boolean	Filter the client request for write with size change.

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.

Name	Type	Description
file_operations	<a href="#">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	<a href="#">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	Specifies the name of the FPolicy event.
protocol	string	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: <ul style="list-style-type: none"> <li>• cifs - for the CIFS protocol.</li> <li>• nfsv3 - for the NFSv3 protocol.</li> <li>• nfsv4 - for the NFSv4 protocol.</li> </ul>
volume_monitoring	boolean	Specifies whether volume operation monitoring is required.

fpolicy\_engine\_reference

FPolicy external engine

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FPolicy external engine.

## fpolicy\_event\_reference

### FPolicy events

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	

### scope

Name	Type	Description
exclude_export_policies	array[string]	
exclude_extension	array[string]	
exclude_shares	array[string]	
exclude_volumes	array[string]	
include_export_policies	array[string]	
include_extension	array[string]	
include_shares	array[string]	
include_volumes	array[string]	

## fpolicy\_policy

Name	Type	Description
enabled	boolean	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	<a href="#">fpolicy_engine_reference</a>	FPolicy external engine
events	array[ <a href="#">fpolicy_event_reference</a> ]	
mandatory	boolean	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	Specifies the name of the policy.

Name	Type	Description
priority	integer	Specifies the priority that is assigned to this policy.
scope	<a href="#">scope</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
engines	array[ <a href="#">fpolicy_engine</a> ]	
events	array[ <a href="#">fpolicy_event</a> ]	
policies	array[ <a href="#">fpolicy_policy</a> ]	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments



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

## Create an FPolicy configuration

POST /protocols/fpolicy

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**

- [DOC /protocols/fpolicy](#)

**Request Body**

Name	Type	Description
_links	<a href="#">_links</a>	
engines	array[ <a href="#">fpolicy_engine</a> ]	
events	array[ <a href="#">fpolicy_event</a> ]	
policies	array[ <a href="#">fpolicy_policy</a> ]	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "_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": "synchronous"
  },
  "events": {
    "name": "event_nfs_close",
    "protocol": "cifs"
  },
  "policies": {
    "engine": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  },
  "events": [
    "event_nfs_close",
    "event_open"
  ],
  "name": "fp_policy_1",
  "scope": {
    "exclude_export_policies": {
    },
    "exclude_extension": {
    },
    "exclude_shares": {
    },
    "exclude_volumes": [

```

```

        "vol1",
        "vol_svm1",
        "*"
    ],
    "include_export_policies": {
    },
    "include_extension": {
    },
    "include_shares": [
        "sh1",
        "share_cifs"
    ],
    "include_volumes": [
        "vol1",
        "vol_svm1"
    ]
    }
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}

```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">fpolicy</a> ]	

## Example response

```
{
  "_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": "synchronous"
    },
    "events": {
      "name": "event_nfs_close",
      "protocol": "cifs"
    },
    "policies": {
      "engine": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      }
    },
    "events": [
      "event_nfs_close",
      "event_open"
    ],
  },
}
```

```

"name": "fp_policy_1",
"scope": {
  "exclude_export_policies": {
  },
  "exclude_extension": {
  },
  "exclude_shares": {
  },
  "exclude_volumes": [
    "voll",
    "vol_svm1",
    "*"
  ],
  "include_export_policies": {
  },
  "include_extension": {
  },
  "include_shares": [
    "sh1",
    "share_cifs"
  ],
  "include_volumes": [
    "voll",
    "vol_svm1"
  ]
},
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"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

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	<a href="#">href</a>	

fpolicy\_engine

The engine defines how ONTAP makes and manages connections to external FPolicy servers.

Name	Type	Description
name	string	Specifies the name to assign to the external server configuration.
port	integer	Port number of the FPolicy server application.
primary_servers	array[string]	
secondary_servers	array[string]	
type	string	The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are: <ul style="list-style-type: none"><li>• synchronous - After sending a notification, wait for a response from the FPolicy server.</li><li>• asynchronous - After sending a notification, file request processing continues.<ul style="list-style-type: none"><li>◦ Default value: 1</li><li>◦ enum: ["synchronous", "asynchronous"]</li></ul></li></ul>

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.



Name	Type	Description
close	boolean	File close operations
create	boolean	File create operations
create_dir	boolean	Directory create operations
delete	boolean	File delete operations
delete_dir	boolean	Directory delete operations
getattr	boolean	Get attribute operations
link	boolean	Link operations
lookup	boolean	Lookup operations
open	boolean	File open operations
read	boolean	File read operations
rename	boolean	File rename operations
rename_dir	boolean	Directory rename operations
setattr	boolean	Set attribute operations
symlink	boolean	Symbolic link operations
write	boolean	File write operations

#### 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.

Name	Type	Description
close_with_modification	boolean	Filter the client request for close with modification.
close_with_read	boolean	Filter the client request for close with read.

<b>Name</b>	<b>Type</b>	<b>Description</b>
close_without_modification	boolean	Filter the client request for close without modification.
exclude_directory	boolean	Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.
first_read	boolean	Filter the client requests for the first-read.
first_write	boolean	Filter the client requests for the first-write.
monitor_ads	boolean	Filter the client request for alternate data stream.
offline_bit	boolean	Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.
open_with_delete_intent	boolean	Filter the client request for open with delete intent.
open_with_write_intent	boolean	Filter the client request for open with write intent.
setattr_with_access_time_change	boolean	Filter the client setattr requests for changing the access time of a file or directory.
setattr_with_allocation_size_change	boolean	Filter the client setattr requests for changing the allocation size of a file.
setattr_with_creation_time_change	boolean	Filter the client setattr requests for changing the creation time of a file or directory.
setattr_with_dacl_change	boolean	Filter the client setattr requests for changing dacl on a file or directory.
setattr_with_group_change	boolean	Filter the client setattr requests for changing group of a file or directory.

Name	Type	Description
setattr_with_mode_change	boolean	Filter the client setattr requests for changing the mode bits on a file or directory.
setattr_with_modify_time_change	boolean	Filter the client setattr requests for changing the modification time of a file or directory.
setattr_with_owner_change	boolean	Filter the client setattr requests for changing owner of a file or directory.
setattr_with_sacl_change	boolean	Filter the client setattr requests for changing sacl on a file or directory.
setattr_with_size_change	boolean	Filter the client setattr requests for changing the size of a file.
write_with_size_change	boolean	Filter the client request for write with size change.

#### 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.

Name	Type	Description
file_operations	<a href="#">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	<a href="#">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	Specifies the name of the FPolicy event.

Name	Type	Description
protocol	string	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: <ul style="list-style-type: none"> <li>• cifs - for the CIFS protocol.</li> <li>• nfsv3 - for the NFSv3 protocol.</li> <li>• nfsv4 - for the NFSv4 protocol.</li> </ul>
volume_monitoring	boolean	Specifies whether volume operation monitoring is required.

fpolicy\_engine\_reference

FPolicy external engine

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FPolicy external engine.

fpolicy\_event\_reference

FPolicy events

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	

scope

Name	Type	Description
exclude_export_policies	array[string]	
exclude_extension	array[string]	
exclude_shares	array[string]	
exclude_volumes	array[string]	
include_export_policies	array[string]	

Name	Type	Description
include_extension	array[string]	
include_shares	array[string]	
include_volumes	array[string]	

#### fpolicy\_policy

Name	Type	Description
enabled	boolean	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	<a href="#">fpolicy_engine_reference</a>	FPolicy external engine
events	array[ <a href="#">fpolicy_event_reference</a> ]	
mandatory	boolean	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	Specifies the name of the policy.
priority	integer	Specifies the priority that is assigned to this policy.
scope	<a href="#">scope</a>	

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### 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.

Name	Type	Description
_links	<a href="#">_links</a>	
engines	array[ <a href="#">fpolicy_engine</a> ]	
events	array[ <a href="#">fpolicy_event</a> ]	
policies	array[ <a href="#">fpolicy_policy</a> ]	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

[\\_links](#)

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete the FPolicy configuration for an SVM

```
DELETE /protocols/fpolicy/{svm.uuid}
```

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

- [DOC /protocols/fpolicy](#)

#### Parameters

Name	Type	In	Required	Description
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
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.

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve the FPolicy configuration for an SVM

GET /protocols/fpolicy/{svm.uuid}

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

- [DOC /protocols/fpolicy](#)

## 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
<a href="#">_links</a>	<a href="#">_links</a>	
engines	array[ <a href="#">fpolicy_engine</a> ]	
events	array[ <a href="#">fpolicy_event</a> ]	
policies	array[ <a href="#">fpolicy_policy</a> ]	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example response

```
{
  "_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": "synchronous"
  },
  "events": {
    "name": "event_nfs_close",
    "protocol": "cifs"
  },
  "policies": {
    "engine": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  },
  "events": [
    "event_nfs_close",
    "event_open"
  ],
  "name": "fp_policy_1",
  "scope": {
    "exclude_export_policies": {
    },
    "exclude_extension": {
    },
    "exclude_shares": {
    },
    "exclude_volumes": [

```

```

        "vol1",
        "vol_svm1",
        "*"
    ],
    "include_export_policies": {
    },
    "include_extension": {
    },
    "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

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	<a href="#">href</a>	

fpolicy\_engine

The engine defines how ONTAP makes and manages connections to external FPolicy servers.

Name	Type	Description
name	string	Specifies the name to assign to the external server configuration.
port	integer	Port number of the FPolicy server application.
primary_servers	array[string]	
secondary_servers	array[string]	
type	string	<p>The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:</p> <ul style="list-style-type: none"><li>• synchronous - After sending a notification, wait for a response from the FPolicy server.</li><li>• asynchronous - After sending a notification, file request processing continues.<ul style="list-style-type: none"><li>◦ Default value: 1</li><li>◦ enum: ["synchronous", "asynchronous"]</li></ul></li></ul>

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.

Name	Type	Description
close	boolean	File close operations
create	boolean	File create operations
create_dir	boolean	Directory create operations
delete	boolean	File delete operations
delete_dir	boolean	Directory delete operations
getattr	boolean	Get attribute operations
link	boolean	Link operations
lookup	boolean	Lookup operations
open	boolean	File open operations
read	boolean	File read operations
rename	boolean	File rename operations
rename_dir	boolean	Directory rename operations
setattr	boolean	Set attribute operations
symlink	boolean	Symbolic link operations
write	boolean	File write operations

#### 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.

Name	Type	Description
close_with_modification	boolean	Filter the client request for close with modification.
close_with_read	boolean	Filter the client request for close with read.

<b>Name</b>	<b>Type</b>	<b>Description</b>
close_without_modification	boolean	Filter the client request for close without modification.
exclude_directory	boolean	Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.
first_read	boolean	Filter the client requests for the first-read.
first_write	boolean	Filter the client requests for the first-write.
monitor_ads	boolean	Filter the client request for alternate data stream.
offline_bit	boolean	Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.
open_with_delete_intent	boolean	Filter the client request for open with delete intent.
open_with_write_intent	boolean	Filter the client request for open with write intent.
setattr_with_access_time_change	boolean	Filter the client setattr requests for changing the access time of a file or directory.
setattr_with_allocation_size_change	boolean	Filter the client setattr requests for changing the allocation size of a file.
setattr_with_creation_time_change	boolean	Filter the client setattr requests for changing the creation time of a file or directory.
setattr_with_dacl_change	boolean	Filter the client setattr requests for changing dacl on a file or directory.
setattr_with_group_change	boolean	Filter the client setattr requests for changing group of a file or directory.

Name	Type	Description
setattr_with_mode_change	boolean	Filter the client setattr requests for changing the mode bits on a file or directory.
setattr_with_modify_time_change	boolean	Filter the client setattr requests for changing the modification time of a file or directory.
setattr_with_owner_change	boolean	Filter the client setattr requests for changing owner of a file or directory.
setattr_with_sacl_change	boolean	Filter the client setattr requests for changing sacl on a file or directory.
setattr_with_size_change	boolean	Filter the client setattr requests for changing the size of a file.
write_with_size_change	boolean	Filter the client request for write with size change.

#### 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.

Name	Type	Description
file_operations	<a href="#">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	<a href="#">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	Specifies the name of the FPolicy event.



Name	Type	Description
protocol	string	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: <ul style="list-style-type: none"> <li>• cifs - for the CIFS protocol.</li> <li>• nfsv3 - for the NFSv3 protocol.</li> <li>• nfsv4 - for the NFSv4 protocol.</li> </ul>
volume_monitoring	boolean	Specifies whether volume operation monitoring is required.

fpolicy\_engine\_reference

FPolicy external engine

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FPolicy external engine.

fpolicy\_event\_reference

FPolicy events

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	

scope

Name	Type	Description
exclude_export_policies	array[string]	
exclude_extension	array[string]	
exclude_shares	array[string]	
exclude_volumes	array[string]	
include_export_policies	array[string]	

Name	Type	Description
include_extension	array[string]	
include_shares	array[string]	
include_volumes	array[string]	

#### fpolicy\_policy

Name	Type	Description
enabled	boolean	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	<a href="#">fpolicy_engine_reference</a>	FPolicy external engine
events	array[ <a href="#">fpolicy_event_reference</a> ]	
mandatory	boolean	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	Specifies the name of the policy.
priority	integer	Specifies the priority that is assigned to this policy.
scope	<a href="#">scope</a>	

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage FPolicy engine configuration

### 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"
```

## Retrieve the FPolicy engine configuration for all engines of an SVM

```
GET /protocols/fpolicy/{svm.uuid}/engines
```

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

- [DOC /protocols/fpolicy/{svm.uuid}/engines](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.



Name	Type	In	Required	Description
primary_servers	string	query	False	Filter by primary_servers
port	integer	query	False	Filter by port
type	string	query	False	Filter by type
secondary_servers	string	query	False	Filter by secondary_servers
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	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">fpolicy_engine</a> ]	

### Example response

```
{
  "_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": "synchronous"
  }
}
```

### Error

Status: Default,

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

fpolicy\_engine

The engine defines how ONTAP makes and manages connections to external FPolicy servers.

Name	Type	Description
name	string	Specifies the name to assign to the external server configuration.
port	integer	Port number of the FPolicy server application.
primary_servers	array[string]	
secondary_servers	array[string]	
type	string	The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are: <ul style="list-style-type: none"><li>• synchronous - After sending a notification, wait for a response from the FPolicy server.</li><li>• asynchronous - After sending a notification, file request processing continues.<ul style="list-style-type: none"><li>◦ Default value: 1</li><li>◦ enum: ["synchronous", "asynchronous"]</li></ul></li></ul>

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Create the FPolicy engine configuration for an SVM

POST /protocols/fpolicy/{svm.uuid}/engines

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

- [DOC /protocols/fpolicy/{svm.uuid}/engines](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

## Request Body

Name	Type	Description
name	string	Specifies the name to assign to the external server configuration.
port	integer	Port number of the FPolicy server application.
primary_servers	array[string]	
secondary_servers	array[string]	
type	string	<p>The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:</p> <ul style="list-style-type: none"><li>• synchronous - After sending a notification, wait for a response from the FPolicy server.</li><li>• asynchronous - After sending a notification, file request processing continues.<ul style="list-style-type: none"><li>◦ Default value: 1</li><li>◦ enum: ["synchronous", "asynchronous"]</li></ul></li></ul>

## Example request

```
{
  "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": "synchronous"
}
```

## Response

Status: 201, Created

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">fpolicy_engine</a> ]	

## Example response

```
{
  "_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": "synchronous"
  }
}
```

## 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

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

### fpolicy\_engine

The engine defines how ONTAP makes and manages connections to external FPolicy servers.

Name	Type	Description
name	string	Specifies the name to assign to the external server configuration.
port	integer	Port number of the FPolicy server application.
primary_servers	array[string]	
secondary_servers	array[string]	
type	string	<p>The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:</p> <ul style="list-style-type: none"><li>• synchronous - After sending a notification, wait for a response from the FPolicy server.</li><li>• asynchronous - After sending a notification, file request processing continues.<ul style="list-style-type: none"><li>◦ Default value: 1</li><li>◦ enum: ["synchronous", "asynchronous"]</li></ul></li></ul>

### href

Name	Type	Description
href	string	

### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete an FPolicy external engine configuration

```
DELETE /protocols/fpolicy/{svm.uuid}/engines/{name}
```

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

- [DOC /protocols/fpolicy/{svm.uuid}/engines](#)

#### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	

#### 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

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve a particular FPolicy engine configuration for an SVM

GET /protocols/fpolicy/{svm.uuid}/engines/{name}

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

- [DOC /protocols/fpolicy/{svm.uuid}/engines](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	

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

### Response

Status: 200, Ok

Name	Type	Description
name	string	Specifies the name to assign to the external server configuration.
port	integer	Port number of the FPolicy server application.
primary_servers	array[string]	
secondary_servers	array[string]	
type	string	<p>The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:</p> <ul style="list-style-type: none"> <li>• synchronous - After sending a notification, wait for a response from the FPolicy server.</li> <li>• asynchronous - After sending a notification, file request processing continues. <ul style="list-style-type: none"> <li>◦ Default value: 1</li> <li>◦ enum: ["synchronous", "asynchronous"]</li> </ul> </li> </ul>

## Example response

```
{
  "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": "synchronous"
}
```

## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update a specific FPolicy engine configuration for an SVM

PATCH /protocols/fpolicy/{svm.uuid}/engines/{name}

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

- [DOC /protocols/fpolicy/{svm.uuid}/engines](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	



## Request Body

Name	Type	Description
name	string	Specifies the name to assign to the external server configuration.
port	integer	Port number of the FPolicy server application.
primary_servers	array[string]	
secondary_servers	array[string]	
type	string	<p>The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:</p> <ul style="list-style-type: none"><li>• synchronous - After sending a notification, wait for a response from the FPolicy server.</li><li>• asynchronous - After sending a notification, file request processing continues.<ul style="list-style-type: none"><li>◦ Default value: 1</li><li>◦ enum: ["synchronous", "asynchronous"]</li></ul></li></ul>

## Example request

```
{
  "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": "synchronous"
}
```

## 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

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

### fpolicy\_engine

The engine defines how ONTAP makes and manages connections to external FPolicy servers.

Name	Type	Description
name	string	Specifies the name to assign to the external server configuration.
port	integer	Port number of the FPolicy server application.
primary_servers	array[string]	
secondary_servers	array[string]	
type	string	<p>The notification mode determines what ONTAP does after sending notifications to FPolicy servers. The possible values are:</p> <ul style="list-style-type: none"><li>• synchronous - After sending a notification, wait for a response from the FPolicy server.</li><li>• asynchronous - After sending a notification, file request processing continues.<ul style="list-style-type: none"><li>◦ Default value: 1</li><li>◦ enum: ["synchronous", "asynchronous"]</li></ul></li></ul>

### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code

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

## Manage FPolicy event configuration

### 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/eventsreturn_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/eventsreturn_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,
        "getattr": 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
}
},
],
"num_records": 2
}
```

## 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"
```

## Retrieve an FPolicy event configuration for all events for an SVM

```
GET /protocols/fpolicy/{svm.uuid}/events
```

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

- [DOC /protocols/fpolicy/{svm.uuid}/events](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	query	False	Filter by name
filters.first_write	boolean	query	False	Filter by filters.first_write
filters.setattr_with_size_change	boolean	query	False	Filter by filters.setattr_with_size_change
filters.monitor_ads	boolean	query	False	Filter by filters.monitor_ads
filters.close_with_read	boolean	query	False	Filter by filters.close_with_read
filters.setattr_with_group_change	boolean	query	False	Filter by filters.setattr_with_group_change
filters.offline_bit	boolean	query	False	Filter by filters.offline_bit
filters.setattr_with_sacl_change	boolean	query	False	Filter by filters.setattr_with_sacl_change
filters.setattr_with_dacl_change	boolean	query	False	Filter by filters.setattr_with_dacl_change
filters.open_with_write_intent	boolean	query	False	Filter by filters.open_with_write_intent
filters.setattr_with_modify_time_change	boolean	query	False	Filter by filters.setattr_with_modify_time_change

Name	Type	In	Required	Description
filters.setattr_with_creation_time_change	boolean	query	False	Filter by filters.setattr_with_creation_time_change
filters.setattr_with_access_time_change	boolean	query	False	Filter by filters.setattr_with_access_time_change
filters.open_with_delete_intent	boolean	query	False	Filter by filters.open_with_delete_intent
filters.setattr_with_allocation_size_change	boolean	query	False	Filter by filters.setattr_with_allocation_size_change
filters.close_without_modification	boolean	query	False	Filter by filters.close_without_modification
filters.write_with_size_change	boolean	query	False	Filter by filters.write_with_size_change
filters.close_with_modification	boolean	query	False	Filter by filters.close_with_modification
filters.exclude_directory	boolean	query	False	Filter by filters.exclude_directory
filters.setattr_with_mode_change	boolean	query	False	Filter by filters.setattr_with_mode_change
filters.first_read	boolean	query	False	Filter by filters.first_read
filters.setattr_with_owner_change	boolean	query	False	Filter by filters.setattr_with_owner_change
protocol	string	query	False	Filter by protocol

<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
volume_monitoring	boolean	query	False	Filter by volume_monitoring
file_operations.link	boolean	query	False	Filter by file_operations.link
file_operations.create	boolean	query	False	Filter by file_operations.create
file_operations.close	boolean	query	False	Filter by file_operations.close
file_operations.setattr	boolean	query	False	Filter by file_operations.setattr
file_operations.rename	boolean	query	False	Filter by file_operations.rename
file_operations.delete	boolean	query	False	Filter by file_operations.delete
file_operations.read	boolean	query	False	Filter by file_operations.read
file_operations.lookup	boolean	query	False	Filter by file_operations.lookup
file_operations.getattr	boolean	query	False	Filter by file_operations.getattr
file_operations.create_dir	boolean	query	False	Filter by file_operations.create_dir
file_operations.rename_dir	boolean	query	False	Filter by file_operations.rename_dir
file_operations.open	boolean	query	False	Filter by file_operations.open

Name	Type	In	Required	Description
file_operations.delete_dir	boolean	query	False	Filter by file_operations.delete_dir
file_operations.write	boolean	query	False	Filter by file_operations.write
file_operations.symmlink	boolean	query	False	Filter by file_operations.symmlink
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">fpolicy_event</a> ]	

### Example response

```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "name": "event_nfs_close",
    "protocol": "cifs"
  }
}

```

### Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	



## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

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.

Name	Type	Description
close	boolean	File close operations
create	boolean	File create operations
create_dir	boolean	Directory create operations
delete	boolean	File delete operations
delete_dir	boolean	Directory delete operations
getattr	boolean	Get attribute operations
link	boolean	Link operations
lookup	boolean	Lookup operations
open	boolean	File open operations
read	boolean	File read operations
rename	boolean	File rename operations
rename_dir	boolean	Directory rename operations
setattr	boolean	Set attribute operations

Name	Type	Description
symlink	boolean	Symbolic link operations
write	boolean	File write operations

#### 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.

Name	Type	Description
close_with_modification	boolean	Filter the client request for close with modification.
close_with_read	boolean	Filter the client request for close with read.
close_without_modification	boolean	Filter the client request for close without modification.
exclude_directory	boolean	Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.
first_read	boolean	Filter the client requests for the first-read.
first_write	boolean	Filter the client requests for the first-write.
monitor_ads	boolean	Filter the client request for alternate data stream.
offline_bit	boolean	Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.
open_with_delete_intent	boolean	Filter the client request for open with delete intent.
open_with_write_intent	boolean	Filter the client request for open with write intent.

Name	Type	Description
setattr_with_access_time_change	boolean	Filter the client setattr requests for changing the access time of a file or directory.
setattr_with_allocation_size_change	boolean	Filter the client setattr requests for changing the allocation size of a file.
setattr_with_creation_time_change	boolean	Filter the client setattr requests for changing the creation time of a file or directory.
setattr_with_dacl_change	boolean	Filter the client setattr requests for changing dacl on a file or directory.
setattr_with_group_change	boolean	Filter the client setattr requests for changing group of a file or directory.
setattr_with_mode_change	boolean	Filter the client setattr requests for changing the mode bits on a file or directory.
setattr_with_modify_time_change	boolean	Filter the client setattr requests for changing the modification time of a file or directory.
setattr_with_owner_change	boolean	Filter the client setattr requests for changing owner of a file or directory.
setattr_with_sacl_change	boolean	Filter the client setattr requests for changing sacl on a file or directory.
setattr_with_size_change	boolean	Filter the client setattr requests for changing the size of a file.
write_with_size_change	boolean	Filter the client request for write with size change.

#### 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.

Name	Type	Description
file_operations	<a href="#">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	<a href="#">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	Specifies the name of the FPolicy event.
protocol	string	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: <ul style="list-style-type: none"> <li>• cifs - for the CIFS protocol.</li> <li>• nfsv3 - for the NFSv3 protocol.</li> <li>• nfsv4 - for the NFSv4 protocol.</li> </ul>
volume_monitoring	boolean	Specifies whether volume operation monitoring is required.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create the FPolicy event configuration for an SVM

POST /protocols/fpolicy/{svm.uuid}/events

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

- [DOC /protocols/fpolicy/{svm.uuid}/events](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

### Request Body

Name	Type	Description
file_operations	<a href="#">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	<a href="#">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	Specifies the name of the FPolicy event.
protocol	string	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: <ul style="list-style-type: none"> <li>• cifs - for the CIFS protocol.</li> <li>• nfsv3 - for the NFSv3 protocol.</li> <li>• nfsv4 - for the NFSv4 protocol.</li> </ul>
volume_monitoring	boolean	Specifies whether volume operation monitoring is required.

### Example request

```
{
  "name": "event_nfs_close",
  "protocol": "cifs"
}
```

### Response

Status: 201, Created

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">fpolicy_event</a> ]	

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "name": "event_nfs_close",
    "protocol": "cifs"
  }
}
```

### 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

Name	Type	Description
error	<a href="#">error</a>	

### Example error

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

### Definitions

## See Definitions

### 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.

Name	Type	Description
close	boolean	File close operations
create	boolean	File create operations
create_dir	boolean	Directory create operations
delete	boolean	File delete operations
delete_dir	boolean	Directory delete operations
getattr	boolean	Get attribute operations
link	boolean	Link operations
lookup	boolean	Lookup operations
open	boolean	File open operations
read	boolean	File read operations
rename	boolean	File rename operations
rename_dir	boolean	Directory rename operations
setattr	boolean	Set attribute operations
symlink	boolean	Symbolic link operations
write	boolean	File write operations

### 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.

<b>Name</b>	<b>Type</b>	<b>Description</b>
close_with_modification	boolean	Filter the client request for close with modification.
close_with_read	boolean	Filter the client request for close with read.
close_without_modification	boolean	Filter the client request for close without modification.
exclude_directory	boolean	Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.
first_read	boolean	Filter the client requests for the first-read.
first_write	boolean	Filter the client requests for the first-write.
monitor_ads	boolean	Filter the client request for alternate data stream.
offline_bit	boolean	Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.
open_with_delete_intent	boolean	Filter the client request for open with delete intent.
open_with_write_intent	boolean	Filter the client request for open with write intent.
setattr_with_access_time_change	boolean	Filter the client setattr requests for changing the access time of a file or directory.
setattr_with_allocation_size_change	boolean	Filter the client setattr requests for changing the allocation size of a file.
setattr_with_creation_time_change	boolean	Filter the client setattr requests for changing the creation time of a file or directory.

Name	Type	Description
setattr_with_dacl_change	boolean	Filter the client setattr requests for changing dacl on a file or directory.
setattr_with_group_change	boolean	Filter the client setattr requests for changing group of a file or directory.
setattr_with_mode_change	boolean	Filter the client setattr requests for changing the mode bits on a file or directory.
setattr_with_modify_time_change	boolean	Filter the client setattr requests for changing the modification time of a file or directory.
setattr_with_owner_change	boolean	Filter the client setattr requests for changing owner of a file or directory.
setattr_with_sacl_change	boolean	Filter the client setattr requests for changing sacl on a file or directory.
setattr_with_size_change	boolean	Filter the client setattr requests for changing the size of a file.
write_with_size_change	boolean	Filter the client request for write with size change.

### 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.

Name	Type	Description
file_operations	<a href="#">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.

Name	Type	Description
filters	<a href="#">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	Specifies the name of the FPolicy event.
protocol	string	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: <ul style="list-style-type: none"> <li>• cifs - for the CIFS protocol.</li> <li>• nfsv3 - for the NFSv3 protocol.</li> <li>• nfsv4 - for the NFSv4 protocol.</li> </ul>
volume_monitoring	boolean	Specifies whether volume operation monitoring is required.

#### href

Name	Type	Description
href	string	

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete a specific FPolicy event configuration for an SVM

```
DELETE /protocols/fpolicy/{svm.uuid}/events/{name}
```

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

- [DOC /protocols/fpolicy/{svm.uuid}/events](#)

#### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	

#### 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

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve a specific FPolicy event configuration for an SVM

GET /protocols/fpolicy/{svm.uuid}/events/{name}

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

- [DOC /protocols/fpolicy/{svm.uuid}/events](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	



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

## Response

Status: 200, Ok

Name	Type	Description
file_operations	<a href="#">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	<a href="#">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	Specifies the name of the FPolicy event.
protocol	string	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: <ul style="list-style-type: none"> <li>• cifs - for the CIFS protocol.</li> <li>• nfsv3 - for the NFSv3 protocol.</li> <li>• nfsv4 - for the NFSv4 protocol.</li> </ul>
volume_monitoring	boolean	Specifies whether volume operation monitoring is required.

## Example response

```
{
  "name": "event_nfs_close",
  "protocol": "cifs"
}
```

## 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

### 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.

Name	Type	Description
close	boolean	File close operations
create	boolean	File create operations
create_dir	boolean	Directory create operations
delete	boolean	File delete operations
delete_dir	boolean	Directory delete operations
getattr	boolean	Get attribute operations
link	boolean	Link operations
lookup	boolean	Lookup operations
open	boolean	File open operations
read	boolean	File read operations
rename	boolean	File rename operations
rename_dir	boolean	Directory rename operations
setattr	boolean	Set attribute operations
symlink	boolean	Symbolic link operations
write	boolean	File write operations

### 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.

<b>Name</b>	<b>Type</b>	<b>Description</b>
close_with_modification	boolean	Filter the client request for close with modification.
close_with_read	boolean	Filter the client request for close with read.
close_without_modification	boolean	Filter the client request for close without modification.
exclude_directory	boolean	Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.
first_read	boolean	Filter the client requests for the first-read.
first_write	boolean	Filter the client requests for the first-write.
monitor_ads	boolean	Filter the client request for alternate data stream.
offline_bit	boolean	Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.
open_with_delete_intent	boolean	Filter the client request for open with delete intent.
open_with_write_intent	boolean	Filter the client request for open with write intent.
setattr_with_access_time_change	boolean	Filter the client setattr requests for changing the access time of a file or directory.
setattr_with_allocation_size_change	boolean	Filter the client setattr requests for changing the allocation size of a file.
setattr_with_creation_time_change	boolean	Filter the client setattr requests for changing the creation time of a file or directory.

Name	Type	Description
setattr_with_dacl_change	boolean	Filter the client setattr requests for changing dacl on a file or directory.
setattr_with_group_change	boolean	Filter the client setattr requests for changing group of a file or directory.
setattr_with_mode_change	boolean	Filter the client setattr requests for changing the mode bits on a file or directory.
setattr_with_modify_time_change	boolean	Filter the client setattr requests for changing the modification time of a file or directory.
setattr_with_owner_change	boolean	Filter the client setattr requests for changing owner of a file or directory.
setattr_with_sacl_change	boolean	Filter the client setattr requests for changing sacl on a file or directory.
setattr_with_size_change	boolean	Filter the client setattr requests for changing the size of a file.
write_with_size_change	boolean	Filter the client request for write with size change.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code

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

## Update a specific FPolicy event configuration for an SVM

PATCH /protocols/fpolicy/{svm.uuid}/events/{name}

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

- [DOC /protocols/fpolicy/{svm.uuid}/events](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	

### Request Body

Name	Type	Description
file_operations	<a href="#">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.

Name	Type	Description
filters	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.
name	string	Specifies the name of the FPolicy event.
protocol	string	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: <ul style="list-style-type: none"> <li>• cifs - for the CIFS protocol.</li> <li>• nfsv3 - for the NFSv3 protocol.</li> <li>• nfsv4 - for the NFSv4 protocol.</li> </ul>
volume_monitoring	boolean	Specifies whether volume operation monitoring is required.

### Example request

```
{
  "name": "event_nfs_close",
  "protocol": "cifs"
}
```

### Response

Status: 200, Ok

### Error

Status: Default

Error Code	Description
9764873	The event is a cluster event

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 file operation or a file operation and filter pair

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

### 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.

Name	Type	Description
close	boolean	File close operations
create	boolean	File create operations
create_dir	boolean	Directory create operations
delete	boolean	File delete operations
delete_dir	boolean	Directory delete operations
getattr	boolean	Get attribute operations
link	boolean	Link operations
lookup	boolean	Lookup operations
open	boolean	File open operations
read	boolean	File read operations
rename	boolean	File rename operations
rename_dir	boolean	Directory rename operations
setattr	boolean	Set attribute operations
symlink	boolean	Symbolic link operations
write	boolean	File write operations

### 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.

<b>Name</b>	<b>Type</b>	<b>Description</b>
close_with_modification	boolean	Filter the client request for close with modification.
close_with_read	boolean	Filter the client request for close with read.
close_without_modification	boolean	Filter the client request for close without modification.
exclude_directory	boolean	Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.
first_read	boolean	Filter the client requests for the first-read.
first_write	boolean	Filter the client requests for the first-write.
monitor_ads	boolean	Filter the client request for alternate data stream.
offline_bit	boolean	Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.
open_with_delete_intent	boolean	Filter the client request for open with delete intent.
open_with_write_intent	boolean	Filter the client request for open with write intent.
setattr_with_access_time_change	boolean	Filter the client setattr requests for changing the access time of a file or directory.
setattr_with_allocation_size_change	boolean	Filter the client setattr requests for changing the allocation size of a file.
setattr_with_creation_time_change	boolean	Filter the client setattr requests for changing the creation time of a file or directory.

Name	Type	Description
setattr_with_dacl_change	boolean	Filter the client setattr requests for changing dacl on a file or directory.
setattr_with_group_change	boolean	Filter the client setattr requests for changing group of a file or directory.
setattr_with_mode_change	boolean	Filter the client setattr requests for changing the mode bits on a file or directory.
setattr_with_modify_time_change	boolean	Filter the client setattr requests for changing the modification time of a file or directory.
setattr_with_owner_change	boolean	Filter the client setattr requests for changing owner of a file or directory.
setattr_with_sacl_change	boolean	Filter the client setattr requests for changing sacl on a file or directory.
setattr_with_size_change	boolean	Filter the client setattr requests for changing the size of a file.
write_with_size_change	boolean	Filter the client request for write with size change.

#### 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.

Name	Type	Description
file_operations	<a href="#">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.

Name	Type	Description
filters	<a href="#">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	Specifies the name of the FPolicy event.
protocol	string	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: <ul style="list-style-type: none"> <li>• cifs - for the CIFS protocol.</li> <li>• nfsv3 - for the NFSv3 protocol.</li> <li>• nfsv4 - for the NFSv4 protocol.</li> </ul>
volume_monitoring	boolean	Specifies whether volume operation monitoring is required.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

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

## Manage SVM FPolicy configuration

### 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 to 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\": [ \"cifs\", \"nfs\" ], \"mandatory\": true, \"name\": \"FPolicy_policy_0\", \"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",
      "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/polices?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"priority\":
1, \"engine\": { \"name\": \"engine1\" }, \"events\": [ \"cifs\", \"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\": [
\"cifs\", \"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/policis?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_pol10"  
        ],  
        "exclude_export_policies": [  
          "export_pol1"  
        ],  
        "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/policis?engine.name=engine1&fields=*&return_records=true&retu
rn_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_pol10"
    ],
    "exclude_export_policies": [
      "export_pol1"
    ],
    "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\": [ \"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 }"
```

---

### Retrieve the FPolicy configuration for an SVM

```
GET /protocols/fpolicy/{svm.uuid}/policies
```

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

- [DOC /protocols/fpolicy/{svm.uuid}/policies](#)

#### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
events.name	string	query	False	Filter by events.name
priority	integer	query	False	Filter by priority
mandatory	boolean	query	False	Filter by mandatory
engine.name	string	query	False	Filter by engine.name
scope.include_shares	string	query	False	Filter by scope.include_shares
scope.include_export_policies	string	query	False	Filter by scope.include_export_policies
scope.include_volumes	string	query	False	Filter by scope.include_volumes
scope.exclude_export_policies	string	query	False	Filter by scope.exclude_export_policies
scope.include_extension	string	query	False	Filter by scope.include_extension

Name	Type	In	Required	Description
scope.exclude_shares	string	query	False	Filter by scope.exclude_shares
scope.exclude_extension	string	query	False	Filter by scope.exclude_extension
scope.exclude_volumes	string	query	False	Filter by scope.exclude_volumes
name	string	query	False	Filter by name
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	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of Records
records	array[ <a href="#">fpolicy_policy</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "engine": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  },
  "events": [
    "event_nfs_close",
    "event_open"
  ],
  "name": "fp_policy_1",
  "scope": {
    "exclude_export_policies": {
    },
    "exclude_extension": {
    },
    "exclude_shares": {
    },
    "exclude_volumes": [
      "vol1",
      "vol_svm1",
      "*"
    ],
    "include_export_policies": {
    },
    "include_extension": {
    },
    "include_shares": [
      "sh1",
      "share_cifs"
    ],
    "include_volumes": [
      "vol1",
```

```
        "vol_svm1"  
      ]  
    }  
  }  
}
```

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

fpolicy\_engine\_reference

FPolicy external engine

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FPolicy external engine.

fpolicy\_event\_reference

FPolicy events

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	

scope

Name	Type	Description
exclude_export_policies	array[string]	
exclude_extension	array[string]	
exclude_shares	array[string]	
exclude_volumes	array[string]	
include_export_policies	array[string]	

Name	Type	Description
include_extension	array[string]	
include_shares	array[string]	
include_volumes	array[string]	

#### fpolicy\_policy

Name	Type	Description
enabled	boolean	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	<a href="#">fpolicy_engine_reference</a>	FPolicy external engine
events	array[ <a href="#">fpolicy_event_reference</a> ]	
mandatory	boolean	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	Specifies the name of the policy.
priority	integer	Specifies the priority that is assigned to this policy.
scope	<a href="#">scope</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error



Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Create the FPolicy configuration for an SVM

POST /protocols/fpolicy/{svm.uuid}/policies

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

- [DOC /protocols/fpolicy/{svm.uuid}/policies](#)

#### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

#### Request Body

Name	Type	Description
enabled	boolean	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	<a href="#">fpolicy_engine_reference</a>	FPolicy external engine
events	array[ <a href="#">fpolicy_event_reference</a> ]	
mandatory	boolean	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	Specifies the name of the policy.
priority	integer	Specifies the priority that is assigned to this policy.
scope	<a href="#">scope</a>	

## Example request

```
{
  "engine": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "events": [
    "event_nfs_close",
    "event_open"
  ],
  "name": "fp_policy_1",
  "scope": {
    "exclude_export_policies": {
    },
    "exclude_extension": {
    },
    "exclude_shares": {
    },
    "exclude_volumes": [
      "vol1",
      "vol_svm1",
      "*"
    ],
    "include_export_policies": {
    },
    "include_extension": {
    },
    "include_shares": [
      "sh1",
      "share_cifs"
    ],
    "include_volumes": [
      "vol1",
      "vol_svm1"
    ]
  }
}
```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of Records
records	array[ <a href="#">fpolicy_policy</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "engine": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  },
  "events": [
    "event_nfs_close",
    "event_open"
  ],
  "name": "fp_policy_1",
  "scope": {
    "exclude_export_policies": {
    },
    "exclude_extension": {
    },
    "exclude_shares": {
    },
    "exclude_volumes": [
      "vol1",
      "vol_svm1",
      "*"
    ],
    "include_export_policies": {
    },
    "include_extension": {
    },
    "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

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	<a href="#">href</a>	

fpolicy\_engine\_reference

FPolicy external engine

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FPolicy external engine.

fpolicy\_event\_reference

FPolicy events

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	

scope

Name	Type	Description
exclude_export_policies	array[string]	
exclude_extension	array[string]	
exclude_shares	array[string]	
exclude_volumes	array[string]	
include_export_policies	array[string]	
include_extension	array[string]	
include_shares	array[string]	
include_volumes	array[string]	

fpolicy\_policy

Name	Type	Description
enabled	boolean	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	<a href="#">fpolicy_engine_reference</a>	FPolicy external engine
events	array[ <a href="#">fpolicy_event_reference</a> ]	
mandatory	boolean	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	Specifies the name of the policy.
priority	integer	Specifies the priority that is assigned to this policy.
scope	<a href="#">scope</a>	

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments



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

### Delete the FPolicy configuration for an SVM

```
DELETE /protocols/fpolicy/{svm.uuid}/policies/{name}
```

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

- [DOC /protocols/fpolicy/{svm.uuid}/policies](#)

#### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	

#### 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

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve the FPolicy configuration for an SVM

```
GET /protocols/fpolicy/{svm.uuid}/policies/{name}
```

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

- [DOC /protocols/fpolicy/{svm.uuid}/policies](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

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

### Response

Status: 200, Ok

Name	Type	Description
enabled	boolean	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	<a href="#">fpolicy_engine_reference</a>	FPolicy external engine
events	array[ <a href="#">fpolicy_event_reference</a> ]	
mandatory	boolean	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	Specifies the name of the policy.
priority	integer	Specifies the priority that is assigned to this policy.
scope	<a href="#">scope</a>	

## Example response

```
{
  "engine": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "events": [
    "event_nfs_close",
    "event_open"
  ],
  "name": "fp_policy_1",
  "scope": {
    "exclude_export_policies": {
    },
    "exclude_extension": {
    },
    "exclude_shares": {
    },
    "exclude_volumes": [
      "vol1",
      "vol_svm1",
      "*"
    ],
    "include_export_policies": {
    },
    "include_extension": {
    },
    "include_shares": [
      "sh1",
      "share_cifs"
    ],
    "include_volumes": [
      "vol1",
      "vol_svm1"
    ]
  }
}
```

## 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	<a href="#">href</a>	

fpolicy\_engine\_reference

FPolicy external engine

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FPolicy external engine.

fpolicy\_event\_reference

FPolicy events

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	

scope

Name	Type	Description
exclude_export_policies	array[string]	
exclude_extension	array[string]	
exclude_shares	array[string]	
exclude_volumes	array[string]	
include_export_policies	array[string]	
include_extension	array[string]	
include_shares	array[string]	
include_volumes	array[string]	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the FPolicy configuration for an SVM

```
PATCH /protocols/fpolicy/{svm.uuid}/policies/{name}
```

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

- [DOC /protocols/fpolicy/{svm.uuid}/policies](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	



## Request Body

Name	Type	Description
enabled	boolean	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	<a href="#">fpolicy_engine_reference</a>	FPolicy external engine
events	array[ <a href="#">fpolicy_event_reference</a> ]	
mandatory	boolean	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	Specifies the name of the policy.
priority	integer	Specifies the priority that is assigned to this policy.
scope	<a href="#">scope</a>	

## Example request

```
{
  "engine": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "events": [
    "event_nfs_close",
    "event_open"
  ],
  "name": "fp_policy_1",
  "scope": {
    "exclude_export_policies": {
    },
    "exclude_extension": {
    },
    "exclude_shares": {
    },
    "exclude_volumes": [
      "vol1",
      "vol_svm1",
      "*"
    ],
    "include_export_policies": {
    },
    "include_extension": {
    },
    "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

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	<a href="#">href</a>	

fpolicy\_engine\_reference

FPolicy external engine

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FPolicy external engine.

fpolicy\_event\_reference

FPolicy events

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	

scope

Name	Type	Description
exclude_export_policies	array[string]	
exclude_extension	array[string]	
exclude_shares	array[string]	
exclude_volumes	array[string]	
include_export_policies	array[string]	
include_extension	array[string]	
include_shares	array[string]	
include_volumes	array[string]	

fpolicy\_policy

Name	Type	Description
enabled	boolean	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	<a href="#">fpolicy_engine_reference</a>	FPolicy external engine
events	array[ <a href="#">fpolicy_event_reference</a> ]	
mandatory	boolean	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	Specifies the name of the policy.
priority	integer	Specifies the priority that is assigned to this policy.
scope	<a href="#">scope</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

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

## Manage NFS export policies

### 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"
test_patch_export_rule.txt (body) :
{
  "new_index": "10",
  "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"
```

---

## Retrieve export policies

```
GET /protocols/nfs/export-policies
```

Retrieves export policies.

## Related ONTAP commands

- `vserver export-policy show`
- `vserver export-policy rule show`

## Learn more

- [DOC /protocols/nfs/export-policies](#)

## Parameters

Name	Type	In	Required	Description
rules.protocols	string	query	False	Filter by rules.protocols
rules.superuser	string	query	False	Filter by rules.superuser
rules.clients.match	string	query	False	Filter by rules.clients.match
rules.ro_rule	string	query	False	Filter by rules.ro_rule
rules.index	integer	query	False	Filter by rules.index
rules.rw_rule	string	query	False	Filter by rules.rw_rule
rules.anonymous_user	string	query	False	Filter by rules.anonymous_user
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
id	integer	query	False	Filter by id
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.

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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	The number of export policy records
records	array[ <a href="#">export_policy</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 0,
    "rules": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "clients": {
        "match": "0.0.0.0/0"
      },
      "index": 0,
      "protocols": {
      },
      "ro_rule": {
      },
      "rw_rule": {
      },
      "superuser": {
      }
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  }
}
```

```
}  
}
```

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

export\_client

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"> <li>• As a hostname; for instance, host1</li> <li>• As an IPv4 address; for instance, 10.1.12.24</li> <li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li> <li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li> <li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li> <li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li> <li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li> <li>• As a domain name preceded by the . character; for instance, .example.com</li> </ul>

#### export\_rule

Name	Type	Description
_links	<a href="#">_links</a>	
anonymous_user	string	User ID To Which Anonymous Users Are Mapped.
clients	array[ <a href="#">export_client</a> ]	Array of client matches
index	integer	Index of the rule within the export policy.
protocols	array[string]	

Name	Type	Description
ro_rule	array[string]	Authentication flavors that the read-only access rule governs
rw_rule	array[string]	Authentication flavors that the read/write access rule governs
superuser	array[string]	Authentication flavors that the superuser security type governs

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

export\_policy

Name	Type	Description
_links	<a href="#">_links</a>	
id	integer	Export Policy ID
name	string	Export Policy Name
rules	array[ <a href="#">export_rule</a> ]	Rules of the Export Policy.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create an export policy

POST `/protocols/nfs/export-policies`

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

- [DOC /protocols/nfs/export-policies](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>id</code>	integer	Export Policy ID
<code>name</code>	string	Export Policy Name
<code>rules</code>	array[ <a href="#">export_rule</a> ]	Rules of the Export Policy.

Name	Type	Description
svm	svm	SVM, applies only to SVM-scoped objects.

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "id": 0,
  "rules": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "clients": {
      "match": "0.0.0.0/0"
    },
    "index": 0,
    "protocols": {
    },
    "ro_rule": {
    },
    "rw_rule": {
    },
    "superuser": {
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	The number of export policy records
records	array[ <a href="#">export_policy</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 0,
    "rules": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "clients": {
        "match": "0.0.0.0/0"
      },
      "index": 0,
      "protocols": {
      },
      "ro_rule": {
      },
      "rw_rule": {
      },
      "superuser": {
      }
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "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

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	<a href="#">href</a>	

export\_client

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"><li>• As a hostname; for instance, host1</li><li>• As an IPv4 address; for instance, 10.1.12.24</li><li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li><li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li><li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li><li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li><li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li><li>• As a domain name preceded by the . character; for instance, .example.com</li></ul>

## export\_rule

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
anonymous_user	string	User ID To Which Anonymous Users Are Mapped.
clients	array[ <a href="#">export_client</a> ]	Array of client matches
index	integer	Index of the rule within the export policy.
protocols	array[string]	
ro_rule	array[string]	Authentication flavors that the read-only access rule governs
rw_rule	array[string]	Authentication flavors that the read/write access rule governs
superuser	array[string]	Authentication flavors that the superuser security type governs

## svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

## export\_policy

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
id	integer	Export Policy ID
name	string	Export Policy Name
rules	array[ <a href="#">export_rule</a> ]	Rules of the Export Policy.

Name	Type	Description
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### export\_client

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"> <li>• As a hostname; for instance, host1</li> <li>• As an IPv4 address; for instance, 10.1.12.24</li> <li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li> <li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li> <li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li> <li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li> <li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li> <li>• As a domain name preceded by the . character; for instance, .example.com</li> </ul>

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### export\_client

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"> <li>• As a hostname; for instance, host1</li> <li>• As an IPv4 address; for instance, 10.1.12.24</li> <li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li> <li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li> <li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li> <li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li> <li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li> <li>• As a domain name preceded by the . character; for instance, .example.com</li> </ul>

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments

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

### Delete an export policy

DELETE /protocols/nfs/export-policies/{id}

Deletes an export policy.

#### Related ONTAP commands

- `vserver export-policy delete`

#### Learn more

- [DOC /protocols/nfs/export-policies](#)

#### Parameters

Name	Type	In	Required	Description
id	integer	path	True	

#### 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

Error Code	Description
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

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an export policy

GET /protocols/nfs/export-policies/{id}

Retrieves an export policy.

### Related ONTAP commands

- `vserver export-policy show`
- `vserver export-policy rule show`

### Learn more

- [DOC /protocols/nfs/export-policies](#)

### Parameters

Name	Type	In	Required	Description
id	integer	path	True	Export Policy ID
fields	array[string]	query	False	Specify the fields to return.



## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
id	integer	Export Policy ID
name	string	Export Policy Name
rules	array[ <a href="#">export_rule</a> ]	Rules of the Export Policy.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "id": 0,
  "rules": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "clients": {
      "match": "0.0.0.0/0"
    },
    "index": 0,
    "protocols": {
    },
    "ro_rule": {
    },
    "rw_rule": {
    },
    "superuser": {
    }
  },
  "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	<a href="#">href</a>	

export\_client

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"><li>• As a hostname; for instance, host1</li><li>• As an IPv4 address; for instance, 10.1.12.24</li><li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li><li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li><li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li><li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li><li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li><li>• As a domain name preceded by the . character; for instance, .example.com</li></ul>

## export\_rule

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
anonymous_user	string	User ID To Which Anonymous Users Are Mapped.
clients	array[ <a href="#">export_client</a> ]	Array of client matches
index	integer	Index of the rule within the export policy.
protocols	array[string]	
ro_rule	array[string]	Authentication flavors that the read-only access rule governs
rw_rule	array[string]	Authentication flavors that the read/write access rule governs
superuser	array[string]	Authentication flavors that the superuser security type governs

## svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments

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

## Update export policy properties

PATCH /protocols/nfs/export-policies/{id}

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

- [DOC /protocols/nfs/export-policies](#)

### Parameters

Name	Type	In	Required	Description
id	integer	path	True	Export Policy ID

### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
id	integer	Export Policy ID
name	string	Export Policy Name
rules	array[ <a href="#">export_rule</a> ]	Rules of the Export Policy.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "id": 0,
  "rules": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "clients": {
      "match": "0.0.0.0/0"
    },
    "index": 0,
    "protocols": {
    },
    "ro_rule": {
    },
    "rw_rule": {
    },
    "superuser": {
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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

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	<a href="#">href</a>	

export\_client

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"><li>• As a hostname; for instance, host1</li><li>• As an IPv4 address; for instance, 10.1.12.24</li><li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li><li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li><li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li><li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li><li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li><li>• As a domain name preceded by the . character; for instance, .example.com</li></ul>

## export\_rule

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
anonymous_user	string	User ID To Which Anonymous Users Are Mapped.
clients	array[ <a href="#">export_client</a> ]	Array of client matches
index	integer	Index of the rule within the export policy.
protocols	array[string]	
ro_rule	array[string]	Authentication flavors that the read-only access rule governs
rw_rule	array[string]	Authentication flavors that the read/write access rule governs
superuser	array[string]	Authentication flavors that the superuser security type governs

## svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

## export\_policy

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
id	integer	Export Policy ID
name	string	Export Policy Name
rules	array[ <a href="#">export_rule</a> ]	Rules of the Export Policy.

Name	Type	Description
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve export policy rules

GET /protocols/nfs/export-policies/{policy.id}/rules

Retrieves export policy rules.

### Related ONTAP commands

- `vserver export-policy rule show`

### Learn more

- [DOC /protocols/nfs/export-policies](#)

### Parameters

Name	Type	In	Required	Description
policy.id	integer	path	True	Export Policy ID
protocols	string	query	False	Filter by protocols

Name	Type	In	Required	Description
superuser	string	query	False	Filter by superuser
clients.match	string	query	False	Filter by clients.match
ro_rule	string	query	False	Filter by ro_rule
index	integer	query	False	Filter by index
rw_rule	string	query	False	Filter by rw_rule
anonymous_user	string	query	False	Filter by anonymous_user
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of Export Rule records
records	array[ <a href="#">export_rule</a> ]	

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "clients": {
      "match": "0.0.0.0/0"
    },
    "index": 0,
    "protocols": {
    },
    "ro_rule": {
    },
    "rw_rule": {
    },
    "superuser": {
    }
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

export\_client

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"> <li>• As a hostname; for instance, host1</li> <li>• As an IPv4 address; for instance, 10.1.12.24</li> <li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li> <li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li> <li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li> <li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li> <li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li> <li>• As a domain name preceded by the . character; for instance, .example.com</li> </ul>

#### export\_rule

Name	Type	Description
_links	<a href="#">_links</a>	
anonymous_user	string	User ID To Which Anonymous Users Are Mapped.
clients	array[ <a href="#">export_client</a> ]	Array of client matches
index	integer	Index of the rule within the export policy.
protocols	array[string]	



Name	Type	Description
ro_rule	array[string]	Authentication flavors that the read-only access rule governs
rw_rule	array[string]	Authentication flavors that the read/write access rule governs
superuser	array[string]	Authentication flavors that the superuser security type governs

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create an export policy rule

POST /protocols/nfs/export-policies/{policy.id}/rules

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

- [DOC /protocols/nfs/export-policies](#)

## Parameters

Name	Type	In	Required	Description
<code>policy.id</code>	integer	path	True	Export Policy ID

## Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>anonymous_user</code>	string	User ID To Which Anonymous Users Are Mapped.
<code>clients</code>	array[ <a href="#">export_client</a> ]	Array of client matches
<code>index</code>	integer	Index of the rule within the export policy.
<code>protocols</code>	array[string]	
<code>ro_rule</code>	array[string]	Authentication flavors that the read-only access rule governs
<code>rw_rule</code>	array[string]	Authentication flavors that the read/write access rule governs
<code>superuser</code>	array[string]	Authentication flavors that the superuser security type governs

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "clients": {
    "match": "0.0.0.0/0"
  },
  "index": 0,
  "protocols": {
  },
  "ro_rule": {
  },
  "rw_rule": {
  },
  "superuser": {
  }
}
```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of Export Rule records
records	array[ <a href="#">export_rule</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "clients": {
      "match": "0.0.0.0/0"
    },
    "index": 0,
    "protocols": {
    },
    "ro_rule": {
    },
    "rw_rule": {
    },
    "superuser": {
    }
  }
}
```

## 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

Error Code	Description
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

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	<a href="#">href</a>	

export\_client

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"><li>• As a hostname; for instance, host1</li><li>• As an IPv4 address; for instance, 10.1.12.24</li><li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li><li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li><li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li><li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li><li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li><li>• As a domain name preceded by the . character; for instance, .example.com</li></ul>

## export\_rule

Name	Type	Description
_links	<a href="#">_links</a>	
anonymous_user	string	User ID To Which Anonymous Users Are Mapped.
clients	array[ <a href="#">export_client</a> ]	Array of client matches
index	integer	Index of the rule within the export policy.
protocols	array[string]	
ro_rule	array[string]	Authentication flavors that the read-only access rule governs
rw_rule	array[string]	Authentication flavors that the read/write access rule governs
superuser	array[string]	Authentication flavors that the superuser security type governs

## \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

## error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

## error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message



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

### Delete an export policy rule

DELETE /protocols/nfs/export-policies/{policy.id}/rules/{index}

Deletes an export policy rule.

### Related ONTAP commands

- `vserver export-policy rule delete`

### Learn more

- [DOC /protocols/nfs/export-policies](#)

### Parameters

Name	Type	In	Required	Description
policy.id	integer	path	True	
index	integer	path	True	

### 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

Name	Type	Description
error	<a href="#">error</a>	

### 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an export policy rule

GET /protocols/nfs/export-policies/{policy.id}/rules/{index}

Retrieves an export policy rule

### Related ONTAP commands

- `vserver export-policy rule show`

### Learn more

- [DOC /protocols/nfs/export-policies](#)

### Parameters

Name	Type	In	Required	Description
policy.id	integer	path	True	Export Policy ID
index	integer	path	True	Export Rule Index
fields	array[string]	query	False	Specify the fields to return.

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
anonymous_user	string	User ID To Which Anonymous Users Are Mapped.
clients	array[ <a href="#">export_client</a> ]	Array of client matches
index	integer	Index of the rule within the export policy.
protocols	array[string]	
ro_rule	array[string]	Authentication flavors that the read-only access rule governs
rw_rule	array[string]	Authentication flavors that the read/write access rule governs

Name	Type	Description
superuser	array[string]	Authentication flavors that the superuser security type governs

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "clients": {
    "match": "0.0.0.0/0"
  },
  "index": 0,
  "protocols": {
  },
  "ro_rule": {
  },
  "rw_rule": {
  },
  "superuser": {
  }
}
```

### 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	<a href="#">href</a>	

export\_client

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"><li>• As a hostname; for instance, host1</li><li>• As an IPv4 address; for instance, 10.1.12.24</li><li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li><li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li><li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li><li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li><li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li><li>• As a domain name preceded by the . character; for instance, .example.com</li></ul>

## error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

## error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the properties of an export policy rule

```
PATCH /protocols/nfs/export-policies/{policy.id}/rules/{index}
```

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

- [DOC /protocols/nfs/export-policies](#)

### Parameters

Name	Type	In	Required	Description
policy.id	integer	path	True	Export Policy ID
index	integer	path	True	Export Rule Index
new_index	integer	query	False	New Export Rule Index

## Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
anonymous_user	string	User ID To Which Anonymous Users Are Mapped.
clients	array[ <a href="#">export_client</a> ]	Array of client matches
index	integer	Index of the rule within the export policy.
protocols	array[string]	
ro_rule	array[string]	Authentication flavors that the read-only access rule governs
rw_rule	array[string]	Authentication flavors that the read/write access rule governs
superuser	array[string]	Authentication flavors that the superuser security type governs

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "clients": {
    "match": "0.0.0.0/0"
  },
  "index": 0,
  "protocols": {
  },
  "ro_rule": {
  },
  "rw_rule": {
  },
  "superuser": {
  }
}
```



## 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

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

export\_client

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"><li>• As a hostname; for instance, host1</li><li>• As an IPv4 address; for instance, 10.1.12.24</li><li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li><li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li><li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li><li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li><li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li><li>• As a domain name preceded by the . character; for instance, .example.com</li></ul>

## export\_rule

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
anonymous_user	string	User ID To Which Anonymous Users Are Mapped.
clients	array[ <a href="#">export_client</a> ]	Array of client matches
index	integer	Index of the rule within the export policy.
protocols	array[string]	
ro_rule	array[string]	Authentication flavors that the read-only access rule governs
rw_rule	array[string]	Authentication flavors that the read/write access rule governs
superuser	array[string]	Authentication flavors that the superuser security type governs

## error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

## error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve export policy rule clients

GET /protocols/nfs/export-policies/{policy.id}/rules/{index}/clients

Retrieves export policy rule clients.

### Learn more

- [DOC /protocols/nfs/export-policies](#)

### Parameters

Name	Type	In	Required	Description
policy.id	integer	path	True	
index	integer	path	True	

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of export rule client records
records	array[ <a href="#">export_client</a> ]	

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "match": "0.0.0.0/0"
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

export\_client

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"> <li>• As a hostname; for instance, host1</li> <li>• As an IPv4 address; for instance, 10.1.12.24</li> <li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li> <li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li> <li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li> <li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li> <li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li> <li>• As a domain name preceded by the . character; for instance, .example.com</li> </ul>

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments



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

### Create an export policy rule client

POST /protocols/nfs/export-policies/{policy.id}/rules/{index}/clients

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

- [DOC /protocols/nfs/export-policies](#)

#### Parameters

Name	Type	In	Required	Description
policy.id	integer	path	True	Export Policy ID
index	integer	path	True	Export Rule Index

#### Request Body

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"><li>• As a hostname; for instance, host1</li><li>• As an IPv4 address; for instance, 10.1.12.24</li><li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li><li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li><li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li><li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li><li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li><li>• As a domain name preceded by the . character; for instance, .example.com</li></ul>

### Example request

```
{  
  "match": "0.0.0.0/0"  
}
```

### Response

```
Status: 201, Created
```

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of export rule client records
records	array[ <a href="#">export_client</a> ]	

### Example response

```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "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

Error Code	Description
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

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

export\_client

Name	Type	Description
match	string	<p>Client Match Hostname, IP Address, Netgroup, or Domain. You can specify the match as a string value in any of the following formats:</p> <ul style="list-style-type: none"><li>• As a hostname; for instance, host1</li><li>• As an IPv4 address; for instance, 10.1.12.24</li><li>• As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1</li><li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li><li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li><li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li><li>• As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng</li><li>• As a domain name preceded by the . character; for instance, .example.com</li></ul>

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
next	<a href="#">href</a>	

Name	Type	Description
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete an export policy client

DELETE /protocols/nfs/export-policies/{policy.id}/rules/{index}/clients/{match}

Deletes an export policy client

#### Related ONTAP commands

- `vserver export-policy rule remove-clientmatches`

#### Learn more

- [DOC /protocols/nfs/export-policies](#)

#### Parameters

Name	Type	In	Required	Description
policy.id	integer	path	True	
index	integer	path	True	
match	string	path	True	

## 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

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## View and update Kerberos interfaces

### 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/{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-vsimg3-d1.sim.netapp.com@NFS-NSR-W01.RTP.NETAPP.COM",
  "user" : "administrator",
  "password" : "Hello123!"
}
```

### Retrieve Kerberos interfaces

```
GET /protocols/nfs/kerberos/interfaces
```

Retrieves Kerberos interfaces.

#### Related ONTAP commands

- `vserver nfs kerberos interface show`

#### Learn more

- [DOC /protocols/nfs/kerberos/interfaces](#)

#### Parameters

Name	Type	In	Required	Description
interface.ip.address	string	query	False	Filter by interface.ip.address

Name	Type	In	Required	Description
interface.uuid	string	query	False	Filter by interface.uuid
interface.name	string	query	False	Filter by interface.name
encryption_types	string	query	False	Filter by encryption_types
spn	string	query	False	Filter by spn
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
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">kerberos_interface</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "encryption_types": {
    },
    "interface": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ip": {
        "address": "10.10.10.7"
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

interface

Network interface

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### kerberos\_interface

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Specifies if Kerberos is enabled.
encryption_types	array[string]	
interface	<a href="#">interface</a>	Network interface
keytab_uri	string	Load keytab from URI
organizational_unit	string	Organizational unit
password	string	Account creation password
spn	string	Service principal name. Valid in PATCH.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
user	string	Account creation user name

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code

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

### Retrieve a Kerberos interface

GET /protocols/nfs/kerberos/interfaces/{uuid}

Retrieves a Kerberos interface.

#### Related ONTAP commands

- `vserver nfs kerberos interface show`

#### Learn more

- [DOC /protocols/nfs/kerberos/interfaces](#)

#### Parameters

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

#### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Specifies if Kerberos is enabled.
encryption_types	array[string]	
interface	<a href="#">interface</a>	Network interface
keytab_uri	string	Load keytab from URI
organizational_unit	string	Organizational unit



Name	Type	Description
password	string	Account creation password
spn	string	Service principal name. Valid in PATCH.
svm	svm	SVM, applies only to SVM-scoped objects.
user	string	Account creation user name

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "encryption_types": {
  },
  "interface": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
  },
  "ip": {
    "address": "10.10.10.7"
  },
  "name": "lif1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"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	<a href="#">href</a>	

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

interface

Network interface

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update Kerberos interface properties

PATCH /protocols/nfs/kerberos/interfaces/{uuid}

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

- [DOC /protocols/nfs/kerberos/interfaces](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Network interface UUID

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
enabled	boolean	Specifies if Kerberos is enabled.
encryption_types	array[string]	
interface	<a href="#">interface</a>	Network interface
keytab_uri	string	Load keytab from URI
organizational_unit	string	Organizational unit
password	string	Account creation password
spn	string	Service principal name. Valid in PATCH.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
user	string	Account creation user name

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "encryption_types": {
  },
  "interface": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
  },
  "ip": {
    "address": "10.10.10.7"
  },
  "name": "lif1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
```

## 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.

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	<a href="#">href</a>	

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

interface

Network interface

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

kerberos\_interface

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	Specifies if Kerberos is enabled.
encryption_types	array[string]	
interface	<a href="#">interface</a>	Network interface
keytab_uri	string	Load keytab from URI
organizational_unit	string	Organizational unit
password	string	Account creation password
spn	string	Service principal name. Valid in PATCH.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
user	string	Account creation user name

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage Kerberos realms

### 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"
```

## Retrieve Kerberos realms

```
GET /protocols/nfs/kerberos/realms
```

Retrieves Kerberos realms.

### Related ONTAP commands

- `vserver nfs kerberos realm show`

### Learn more

- [DOC /protocols/nfs/kerberos/realms](#)

## Parameters

Name	Type	In	Required	Description
name	string	query	False	Filter by name
kdc.ip	string	query	False	Filter by kdc.ip
kdc.vendor	string	query	False	Filter by kdc.vendor
kdc.port	integer	query	False	Filter by kdc.port
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
encryption_types	string	query	False	Filter by encryption_types
ad_server.name	string	query	False	Filter by ad_server.name
ad_server.address	string	query	False	Filter by ad_server.address
comment	string	query	False	Filter by comment
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.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">kerberos_realm</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ad_server": {
      "address": "1.2.3.4"
    },
    "comment": "string",
    "encryption_types": {
    },
    "kdc": {
      "ip": "1.2.3.4",
      "port": 88,
      "vendor": "microsoft"
    },
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

ad\_server

Name	Type	Description
address	string	Active Directory server IP address
name	string	Active Directory server name

kdc

Name	Type	Description
ip	string	KDC IP address
port	integer	KDC port
vendor	string	Key Distribution Center (KDC) vendor. Following values are supported: <ul style="list-style-type: none"><li>• microsoft - Microsoft Active Directory KDC</li><li>• other - MIT Kerberos KDC or other KDC</li></ul>

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### kerberos\_realm

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
<a href="#">ad_server</a>	<a href="#">ad_server</a>	
comment	string	Comment
encryption_types	array[string]	
<a href="#">kdc</a>	<a href="#">kdc</a>	
name	string	Kerberos realm
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create a Kerberos realm

POST /protocols/nfs/kerberos/realms

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

- [DOC /protocols/nfs/kerberos/realms](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>ad_server</code>	<a href="#">ad_server</a>	
<code>comment</code>	string	Comment
<code>encryption_types</code>	array[string]	
<code>kdc</code>	<a href="#">kdc</a>	
<code>name</code>	string	Kerberos realm

Name	Type	Description
svm	svm	SVM, applies only to SVM-scoped objects.

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_server": {
    "address": "1.2.3.4"
  },
  "comment": "string",
  "encryption_types": {
  },
  "kdc": {
    "ip": "1.2.3.4",
    "port": 88,
    "vendor": "microsoft"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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.

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	<a href="#">href</a>	

ad\_server

Name	Type	Description
address	string	Active Directory server IP address
name	string	Active Directory server name

kdc

Name	Type	Description
ip	string	KDC IP address
port	integer	KDC port
vendor	string	Key Distribution Center (KDC) vendor. Following values are supported: <ul style="list-style-type: none"><li>• microsoft - Microsoft Active Directory KDC</li><li>• other - MIT Kerberos KDC or other KDC</li></ul>

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.

Name	Type	Description
uuid	string	The unique identifier of the SVM.

#### kerberos\_realm

Name	Type	Description
_links	<a href="#">_links</a>	
ad_server	<a href="#">ad_server</a>	
comment	string	Comment
encryption_types	array[string]	
kdc	<a href="#">kdc</a>	
name	string	Kerberos realm
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete a Kerberos realm

DELETE /protocols/nfs/kerberos/realms/{svm.uuid}/{name}

Deletes a Kerberos realm.

- `vserver nfs kerberos realm delete`

### Learn more

- [DOC /protocols/nfs/kerberos/realms](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	
name	string	path	True	

### 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.

Name	Type	Description
error	<a href="#">error</a>	



## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve a Kerberos realm

GET /protocols/nfs/kerberos/realms/{svm.uuid}/{name}

Retrieves a Kerberos realm.

- `vserver nfs kerberos realm show`

### Learn more

- [DOC /protocols/nfs/kerberos/realms](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	SVM UUID
name	string	path	True	Kerberos realm
fields	array[string]	query	False	Specify the fields to return.

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
<a href="#">ad_server</a>	<a href="#">ad_server</a>	
comment	string	Comment
encryption_types	array[string]	
<a href="#">kdc</a>	<a href="#">kdc</a>	
name	string	Kerberos realm
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_server": {
    "address": "1.2.3.4"
  },
  "comment": "string",
  "encryption_types": {
  },
  "kdc": {
    "ip": "1.2.3.4",
    "port": 88,
    "vendor": "microsoft"
  },
  "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	<a href="#">href</a>	

ad\_server

Name	Type	Description
address	string	Active Directory server IP address
name	string	Active Directory server name

kdc

Name	Type	Description
ip	string	KDC IP address
port	integer	KDC port
vendor	string	Key Distribution Center (KDC) vendor. Following values are supported: <ul style="list-style-type: none"><li>• microsoft - Microsoft Active Directory KDC</li><li>• other - MIT Kerberos KDC or other KDC</li></ul>

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update Kerberos realm properties

PATCH /protocols/nfs/kerberos/realms/{svm.uuid}/{name}

Updates the properties of a Kerberos realm.

- `vserver nfs kerberos realm modify`

### Learn more

- [DOC /protocols/nfs/kerberos/realms](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	SVM UUID
name	string	path	True	Kerberos realm

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
ad_server	<a href="#">ad_server</a>	
comment	string	Comment
encryption_types	array[string]	
kdc	<a href="#">kdc</a>	
name	string	Kerberos realm
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_server": {
    "address": "1.2.3.4"
  },
  "comment": "string",
  "encryption_types": {
  },
  "kdc": {
    "ip": "1.2.3.4",
    "port": 88,
    "vendor": "microsoft"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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.



Error codes	Description
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.

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	<a href="#">href</a>	

ad\_server

Name	Type	Description
address	string	Active Directory server IP address
name	string	Active Directory server name

kdc

Name	Type	Description
ip	string	KDC IP address
port	integer	KDC port
vendor	string	Key Distribution Center (KDC) vendor. Following values are supported: <ul style="list-style-type: none"><li>• microsoft - Microsoft Active Directory KDC</li><li>• other - MIT Kerberos KDC or other KDC</li></ul>

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.

Name	Type	Description
uuid	string	The unique identifier of the SVM.

#### kerberos\_realm

Name	Type	Description
_links	<a href="#">_links</a>	
ad_server	<a href="#">ad_server</a>	
comment	string	Comment
encryption_types	array[string]	
kdc	<a href="#">kdc</a>	
name	string	Kerberos realm
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage NFS services

## 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"
```

### Retrieve NFS configuration for SVMs

GET /protocols/nfs/services

Retrieves the NFS configuration of SVMs.

#### Related ONTAP commands

- `vserver nfs show`
- `vserver nfs status`

#### Learn more

- [DOC /protocols/nfs/services](#)

#### Parameters

Name	Type	In	Required	Description
vstorage_enabled	boolean	query	False	Filter by vstorage_enabled
svm.uuid	string	query	False	Filter by svm.uuid

<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
svm.name	string	query	False	Filter by svm.name
protocol.v41_features.write_delegation_enabled	boolean	query	False	Filter by protocol.v41_features.write_delegation_enabled
protocol.v41_features.acl_enabled	boolean	query	False	Filter by protocol.v41_features.acl_enabled
protocol.v41_features.read_delegation_enabled	boolean	query	False	Filter by protocol.v41_features.read_delegation_enabled
protocol.v41_features.pnfs_enabled	boolean	query	False	Filter by protocol.v41_features.pnfs_enabled
protocol.v40_enabled	boolean	query	False	Filter by protocol.v40_enabled
protocol.v41_enabled	boolean	query	False	Filter by protocol.v41_enabled
protocol.v4_id_domain	string	query	False	Filter by protocol.v4_id_domain
protocol.v40_features.acl_enabled	boolean	query	False	Filter by protocol.v40_features.acl_enabled
protocol.v40_features.write_delegation_enabled	boolean	query	False	Filter by protocol.v40_features.write_delegation_enabled
protocol.v40_features.read_delegation_enabled	boolean	query	False	Filter by protocol.v40_features.read_delegation_enabled

Name	Type	In	Required	Description
protocol.v3_enabled	boolean	query	False	Filter by protocol.v3_enabled
transport.udp_enabled	boolean	query	False	Filter by transport.udp_enabled
transport.tcp_enabled	boolean	query	False	Filter by transport.tcp_enabled
state	string	query	False	Filter by state
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_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.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of NFS Server Records
records	array[ <a href="#">nfs_service</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "state": "online",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

v40\_features

Name	Type	Description
acl_enabled	boolean	Specifies whether NFSv4.0 ACLs is enabled.
read_delegation_enabled	boolean	Specifies whether NFSv4.0 Read Delegation is enabled.
write_delegation_enabled	boolean	Specifies whether NFSv4.0 Write Delegation is enabled.

v41\_features

Name	Type	Description
acl_enabled	boolean	Specifies whether NFSv4.1 ACLs is enabled.
pnfs_enabled	boolean	Specifies whether NFSv4.1 Parallel NFS is enabled.
read_delegation_enabled	boolean	Specifies whether NFSv4.1 Read Delegation is enabled.
write_delegation_enabled	boolean	Specifies whether NFSv4.1 Write Delegation is enabled.

## protocol

Name	Type	Description
v3_enabled	boolean	Specifies whether NFSv3 protocol is enabled.
v40_enabled	boolean	Specifies whether NFSv4.0 protocol is enabled.
v40_features	<a href="#">v40_features</a>	
v41_enabled	boolean	Specifies whether NFSv4.1 protocol is enabled.
v41_features	<a href="#">v41_features</a>	
v4_id_domain	string	Specifies the domain portion of the string form of user and group names as defined by the NFSv4 protocol.

## svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

## transport

Name	Type	Description
tcp_enabled	boolean	Specifies whether TCP transports are enabled on the server.
udp_enabled	boolean	Specifies whether UDP transports are enabled on the server.

## nfs\_service

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
enabled	boolean	Specifies if the NFS service is administratively enabled.
protocol	<a href="#">protocol</a>	
state	string	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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
transport	<a href="#">transport</a>	
vstorage_enabled	boolean	Specifies whether VMware vstorage feature is enabled.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create the NFS configuration for an SVM

POST /protocols/nfs/services

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.v4_id_domain` - *defaultv4iddomain.com*
- `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*

## Related ONTAP commands

- `vserver nfs create`

## Learn more

- [DOC /protocols/nfs/services](#)

## Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>enabled</code>	boolean	Specifies if the NFS service is administratively enabled.
<code>protocol</code>	<a href="#">protocol</a>	

Name	Type	Description
state	string	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	svm	SVM, applies only to SVM-scoped objects.
transport	transport	
vstorage_enabled	boolean	Specifies whether VMware vstorage feature is enabled.

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "state": "online",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
```

### Response

Status: 201, Created

Name	Type	Description
_links	_links	

Name	Type	Description
num_records	integer	Number of NFS Server Records
records	array[nfs_service]	

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "state": "online",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "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

Name	Type	Description
error	<a href="#">error</a>	



## 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	<a href="#">href</a>	

v40\_features

Name	Type	Description
acl_enabled	boolean	Specifies whether NFSv4.0 ACLs is enabled.
read_delegation_enabled	boolean	Specifies whether NFSv4.0 Read Delegation is enabled.
write_delegation_enabled	boolean	Specifies whether NFSv4.0 Write Delegation is enabled.

v41\_features

Name	Type	Description
acl_enabled	boolean	Specifies whether NFSv4.1 ACLs is enabled.
pnfs_enabled	boolean	Specifies whether NFSv4.1 Parallel NFS is enabled.
read_delegation_enabled	boolean	Specifies whether NFSv4.1 Read Delegation is enabled.
write_delegation_enabled	boolean	Specifies whether NFSv4.1 Write Delegation is enabled.

protocol

Name	Type	Description
v3_enabled	boolean	Specifies whether NFSv3 protocol is enabled.

Name	Type	Description
v40_enabled	boolean	Specifies whether NFSv4.0 protocol is enabled.
v40_features	<a href="#">v40_features</a>	
v41_enabled	boolean	Specifies whether NFSv4.1 protocol is enabled.
v41_features	<a href="#">v41_features</a>	
v4_id_domain	string	Specifies the domain portion of the string form of user and group names as defined by the NFSv4 protocol.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

transport

Name	Type	Description
tcp_enabled	boolean	Specifies whether TCP transports are enabled on the server.
udp_enabled	boolean	Specifies whether UDP transports are enabled on the server.

nfs\_service

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Specifies if the NFS service is administratively enabled.
protocol	<a href="#">protocol</a>	

Name	Type	Description
state	string	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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
transport	<a href="#">transport</a>	
vstorage_enabled	boolean	Specifies whether VMware vstorage feature is enabled.

#### [\\_links](#)

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### nfs\_service

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Specifies if the NFS service is administratively enabled.
protocol	<a href="#">protocol</a>	
state	string	Specifies the state of the NFS service on the SVM. The following values are supported: <ul style="list-style-type: none"> <li>• online - NFS server is ready to accept client requests.</li> <li>• offline - NFS server is not ready to accept client requests.</li> </ul>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
transport	<a href="#">transport</a>	

Name	Type	Description
vstorage_enabled	boolean	Specifies whether VMware vstorage feature is enabled.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete the NFS configuration for an SVM

```
DELETE /protocols/nfs/services/{svm.uuid}
```

Deletes the NFS configuration of an SVM.

### Related ONTAP commands

- `vserver nfs delete`

### Learn more

- [DOC /protocols/nfs/services](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	

### 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

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve the NFS configuration for an SVM

GET /protocols/nfs/services/{svm.uuid}

Retrieves the NFS configuration of an SVM.

### Related ONTAP commands

- `vserver nfs show`
- `vserver nfs status`

### Learn more

- [DOC /protocols/nfs/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	<a href="#">_links</a>	
enabled	boolean	Specifies if the NFS service is administratively enabled.
protocol	<a href="#">protocol</a>	
state	string	Specifies the state of the NFS service on the SVM. The following values are supported: <ul style="list-style-type: none"><li>• online - NFS server is ready to accept client requests.</li><li>• offline - NFS server is not ready to accept client requests.</li></ul>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
transport	<a href="#">transport</a>	
vstorage_enabled	boolean	Specifies whether VMware vstorage feature is enabled.



## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "state": "online",
  "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	<a href="#">href</a>	

v40\_features

Name	Type	Description
acl_enabled	boolean	Specifies whether NFSv4.0 ACLs is enabled.
read_delegation_enabled	boolean	Specifies whether NFSv4.0 Read Delegation is enabled.
write_delegation_enabled	boolean	Specifies whether NFSv4.0 Write Delegation is enabled.

v41\_features

Name	Type	Description
acl_enabled	boolean	Specifies whether NFSv4.1 ACLs is enabled.
pnfs_enabled	boolean	Specifies whether NFSv4.1 Parallel NFS is enabled.
read_delegation_enabled	boolean	Specifies whether NFSv4.1 Read Delegation is enabled.
write_delegation_enabled	boolean	Specifies whether NFSv4.1 Write Delegation is enabled.

protocol

Name	Type	Description
v3_enabled	boolean	Specifies whether NFSv3 protocol is enabled.

Name	Type	Description
v40_enabled	boolean	Specifies whether NFSv4.0 protocol is enabled.
v40_features	<a href="#">v40_features</a>	
v41_enabled	boolean	Specifies whether NFSv4.1 protocol is enabled.
v41_features	<a href="#">v41_features</a>	
v4_id_domain	string	Specifies the domain portion of the string form of user and group names as defined by the NFSv4 protocol.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

transport

Name	Type	Description
tcp_enabled	boolean	Specifies whether TCP transports are enabled on the server.
udp_enabled	boolean	Specifies whether UDP transports are enabled on the server.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the NFS configuration for an SVM

PATCH /protocols/nfs/services/{svm.uuid}

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

- [DOC /protocols/nfs/services](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	Specifies if the NFS service is administratively enabled.
protocol	<a href="#">protocol</a>	

Name	Type	Description
state	string	Specifies the state of the NFS service on the SVM. The following values are supported: <ul style="list-style-type: none"> <li>• online - NFS server is ready to accept client requests.</li> <li>• offline - NFS server is not ready to accept client requests.</li> </ul>
svm	svm	SVM, applies only to SVM-scoped objects.
transport	transport	
vstorage_enabled	boolean	Specifies whether VMware vstorage feature is enabled.

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "state": "online",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "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

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

v40\_features

Name	Type	Description
acl_enabled	boolean	Specifies whether NFSv4.0 ACLs is enabled.
read_delegation_enabled	boolean	Specifies whether NFSv4.0 Read Delegation is enabled.
write_delegation_enabled	boolean	Specifies whether NFSv4.0 Write Delegation is enabled.

v41\_features

Name	Type	Description
acl_enabled	boolean	Specifies whether NFSv4.1 ACLs is enabled.
pnfs_enabled	boolean	Specifies whether NFSv4.1 Parallel NFS is enabled.
read_delegation_enabled	boolean	Specifies whether NFSv4.1 Read Delegation is enabled.
write_delegation_enabled	boolean	Specifies whether NFSv4.1 Write Delegation is enabled.

protocol

Name	Type	Description
v3_enabled	boolean	Specifies whether NFSv3 protocol is enabled.

Name	Type	Description
v40_enabled	boolean	Specifies whether NFSv4.0 protocol is enabled.
v40_features	<a href="#">v40_features</a>	
v41_enabled	boolean	Specifies whether NFSv4.1 protocol is enabled.
v41_features	<a href="#">v41_features</a>	
v4_id_domain	string	Specifies the domain portion of the string form of user and group names as defined by the NFSv4 protocol.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

transport

Name	Type	Description
tcp_enabled	boolean	Specifies whether TCP transports are enabled on the server.
udp_enabled	boolean	Specifies whether UDP transports are enabled on the server.

nfs\_service

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	Specifies if the NFS service is administratively enabled.
protocol	<a href="#">protocol</a>	

Name	Type	Description
state	string	Specifies the state of the NFS service on the SVM. The following values are supported: <ul style="list-style-type: none"> <li>• online - NFS server is ready to accept client requests.</li> <li>• offline - NFS server is not ready to accept client requests.</li> </ul>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
transport	<a href="#">transport</a>	
vstorage_enabled	boolean	Specifies whether VMware vstorage feature is enabled.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## View and create Vscan configuration

### Protocols Vscan endpoint overview

#### Overview

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 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*"
      ]
    }
  },

```



```

    "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": [
      "/voll"
    ],
    "log_path": "/voll",
    "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"
```

## Retrieve the Vscan configuration

GET /protocols/vscan

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:

- There can be only one Vscan configuration enabled for an SVM at any time.
- 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

- [DOC /protocols/vscan](#)
- [DOC /protocols/vscan/{svm.uuid}/scanner-pools](#)

### Parameters

Name	Type	In	Required	Description
enabled	boolean	query	False	Filter by enabled

Name	Type	In	Required	Description
on_access_policies.enabled	boolean	query	False	Filter by on_access_policies.enabled
on_access_policies.name	string	query	False	Filter by on_access_policies.name
on_access_policies.mandatory	boolean	query	False	Filter by on_access_policies.mandatory
on_access_policies.scope.include_extensions	string	query	False	Filter by on_access_policies.scope.include_extensions
on_access_policies.scope.scan_readonly_volumes	boolean	query	False	Filter by on_access_policies.scope.scan_readonly_volumes
on_access_policies.scope.exclude_extensions	string	query	False	Filter by on_access_policies.scope.exclude_extensions
on_access_policies.scope.max_file_size	integer	query	False	Filter by on_access_policies.scope.max_file_size
on_access_policies.scope.scan_without_extension	boolean	query	False	Filter by on_access_policies.scope.scan_without_extension
on_access_policies.scope.exclude_paths	string	query	False	Filter by on_access_policies.scope.exclude_paths
on_access_policies.scope.only_execute_access	boolean	query	False	Filter by on_access_policies.scope.only_execute_access

Name	Type	In	Required	Description
scanner_pools.privileged_users	string	query	False	Filter by scanner_pools.privileged_users
scanner_pools.name	string	query	False	Filter by scanner_pools.name
scanner_pools.cluster.uuid	string	query	False	Filter by scanner_pools.cluster.uuid
scanner_pools.cluster.name	string	query	False	Filter by scanner_pools.cluster.name
scanner_pools.servers	string	query	False	Filter by scanner_pools.servers
scanner_pools.role	string	query	False	Filter by scanner_pools.role
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
on_demand_policies.log_path	string	query	False	Filter by on_demand_policies.log_path
on_demand_policies.scope.include_extensions	string	query	False	Filter by on_demand_policies.scope.include_extensions
on_demand_policies.scope.exclude_extensions	string	query	False	Filter by on_demand_policies.scope.exclude_extensions
on_demand_policies.scope.max_file_size	integer	query	False	Filter by on_demand_policies.scope.max_file_size

Name	Type	In	Required	Description
on_demand_policies.scope.exclude_paths	string	query	False	Filter by on_demand_policies.scope.exclude_paths
on_demand_policies.scope.scan_without_extension	boolean	query	False	Filter by on_demand_policies.scope.scan_without_extension
on_demand_policies.schedule.uuid	string	query	False	Filter by on_demand_policies.schedule.uuid
on_demand_policies.schedule.name	string	query	False	Filter by on_demand_policies.schedule.name
on_demand_policies.name	string	query	False	Filter by on_demand_policies.name
on_demand_policies.scan_paths	string	query	False	Filter by on_demand_policies.scan_paths
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.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">vscan</a> ]	

## Example response

```
{
  "_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": [
        "/voll/cold-files/",
        "/voll/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": "primary",
    "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

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
only_execute_access	boolean	Scan only files opened with execute-access.
scan_readonly_volumes	boolean	Specifies whether or not read-only volume can be scanned.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

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.

Name	Type	Description
enabled	boolean	Status of the On-Access Vscan policy
mandatory	boolean	Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning.
name	string	On-Access policy name
scope	<a href="#">scope</a>	

schedule

Schedule of the task.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

## 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.

Name	Type	Description
log_path	string	The path from the Vserver root where the task report is created.
name	string	On-Demand task name
scan_paths	array[string]	List of paths that need to be scanned.
schedule	<a href="#">schedule</a>	Schedule of the task.
scope	<a href="#">scope</a>	

## cluster\_reference

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

## 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".

Name	Type	Description
cluster	<a href="#">cluster_reference</a>	
name	string	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 ".".

Name	Type	Description
privileged_users	array[string]	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. <ul style="list-style-type: none"> <li>example: ["cifs\u1", "cifs\u2"]</li> </ul>
role	string	Specifies the role of the scanner pool. The possible values are: <ul style="list-style-type: none"> <li>primary - Always active.</li> <li>secondary - Active only when none of the primary external virus-scanning servers are connected.</li> <li>idle - Always inactive.</li> </ul>
servers	array[string]	Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP. <ul style="list-style-type: none"> <li>example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]</li> </ul>

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
cache_clear	boolean	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	Specifies whether or not Vscan is enabled on the SVM.
on_access_policies	array[ <a href="#">vscan_on_access</a> ]	
on_demand_policies	array[ <a href="#">vscan_on_demand</a> ]	
scanner_pools	array[ <a href="#">vscan_scanner_pool</a> ]	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create a Vscan configuration

POST /protocols/vscan

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're creating is enabled or disabled for a specified SVM.

Important notes:

- There can be only one Vscan configuration enabled for an SVM at any time.
- There needs to be at least one active scanner-pool and one enabled On-Access policy for Vscan to be enabled 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

- [DOC /protocols/vscan](#)
- [DOC /protocols/vscan/{svm.uuid}/scanner-pools](#)

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
cache_clear	boolean	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	Specifies whether or not Vscan is enabled on the SVM.
on_access_policies	array[ <a href="#">vscan_on_access</a> ]	
on_demand_policies	array[ <a href="#">vscan_on_demand</a> ]	
scanner_pools	array[ <a href="#">vscan_scanner_pool</a> ]	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "_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": [
        "/voll/cold-files/",
        "/voll/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": "primary",
    "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"
}
}

```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">vscan</a> ]	

## Example response

```
{
  "_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": [
        "/voll/cold-files/",
        "/voll/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": "primary",
    "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

## ONTAP Error Response Codes

Error Code	Description
10027259 +	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."
10027260 +	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."
2621462 +	The specified SVM name is invalid
2621706 +	The specified svm.uuid is either invalid or belongs to a different SVM
10027015 +	Attempting to enable a Vscan but no active scanner-pool exists for the specified SVM
10027011 +	Attempting to enable a Vscan for an SVM for which no CIFS server exists
10027023 +	Attempting to enable a Vscan for an SVM for which no active Vscan On-Access policy 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
only_execute_access	boolean	Scan only files opened with execute-access.
scan_readonly_volumes	boolean	Specifies whether or not read-only volume can be scanned.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

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.

Name	Type	Description
enabled	boolean	Status of the On-Access Vscan policy

Name	Type	Description
mandatory	boolean	Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning.
name	string	On-Access policy ame
scope	<a href="#">scope</a>	

schedule

Schedule of the task.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

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.

Name	Type	Description
log_path	string	The path from the Vserver root where the task report is created.
name	string	On-Demand task name
scan_paths	array[string]	List of paths that need to be scanned.
schedule	<a href="#">schedule</a>	Schedule of the task.
scope	<a href="#">scope</a>	

#### cluster\_reference

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

#### 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".

Name	Type	Description
cluster	<a href="#">cluster_reference</a>	
name	string	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 ".".

Name	Type	Description
privileged_users	array[string]	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. <ul style="list-style-type: none"> <li>example: ["cifs\u1", "cifs\u2"]</li> </ul>
role	string	Specifies the role of the scanner pool. The possible values are: <ul style="list-style-type: none"> <li>primary - Always active.</li> <li>secondary - Active only when none of the primary external virus-scanning servers are connected.</li> <li>idle - Always inactive.</li> </ul>
servers	array[string]	Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP. <ul style="list-style-type: none"> <li>example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]</li> </ul>

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
cache_clear	boolean	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	Specifies whether or not Vscan is enabled on the SVM.
on_access_policies	array[ <a href="#">vscan_on_access</a> ]	
on_demand_policies	array[ <a href="#">vscan_on_demand</a> ]	
scanner_pools	array[ <a href="#">vscan_scanner_pool</a> ]	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

[\\_links](#)

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code

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

## Manage Vscan configuration

### 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=*"
    }
  }
}

```

## Retrieve the Vscan server status

GET /protocols/vscan/server-status

Retrieves a Vscan server status.

### Related ONTAP commands

- `vserver vscan connection-status show-all`

### Learn more

- [DOC /protocols/vscan/server-status](#)

### Parameters

Name	Type	In	Required	Description
ip	string	query	False	Filter by ip
disconnected_reason	string	query	False	Filter by disconnected_reason
version	string	query	False	Filter by version

Name	Type	In	Required	Description
update_time	string	query	False	Filter by update_time
state	string	query	False	Filter by state
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
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
vendor	string	query	False	Filter by vendor
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">vscan_server_status</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "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"
    },
    "type": "primary"
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

vscan\_server\_status

Displays the connection status information of the external virus-scanning servers.



Name	Type	Description
disconnected_reason	string	<p data-bbox="1036 155 1453 218">Specifies the server disconnected reason.</p> <p data-bbox="1036 226 1372 289">The following is a list of the possible reasons:</p> <ul data-bbox="1058 323 1453 1428" style="list-style-type: none"> <li data-bbox="1058 323 1396 386">• unknown - Disconnected, unknown reason.</li> <li data-bbox="1058 415 1372 508">• vscan_disabled - Disconnected, Vscan is disabled on the SVM.</li> <li data-bbox="1058 529 1442 592">• no_data_lif - Disconnected, SVM does not have data LIF.</li> <li data-bbox="1058 613 1442 705">• session_uninitialized - Disconnected, session is not initialized.</li> <li data-bbox="1058 726 1396 819">• remote_closed - Disconnected, server has closed the connection.</li> <li data-bbox="1058 840 1453 932">• invalid_protocol_msg - Disconnected, invalid protocol message received.</li> <li data-bbox="1058 953 1453 1045">• invalid_session_id - Disconnected, invalid session ID received.</li> <li data-bbox="1058 1066 1429 1159">• inactive_connection - Disconnected, no activity on connection.</li> <li data-bbox="1058 1180 1404 1272">• invalid_user - Connection request by an invalid user.</li> <li data-bbox="1058 1293 1453 1386">• server_removed - Disconnected, server has been removed from the active Scanners List.</li> </ul> <p data-bbox="1081 1415 1161 1444">enum:</p> <ul data-bbox="1058 1465 1339 1957" style="list-style-type: none"> <li data-bbox="1058 1465 1193 1495">• unknown</li> <li data-bbox="1058 1516 1274 1545">• vscan_disabled</li> <li data-bbox="1058 1566 1226 1596">• no_data_lif</li> <li data-bbox="1058 1617 1339 1646">• session_uninitialized</li> <li data-bbox="1058 1667 1266 1696">• remote_closed</li> <li data-bbox="1058 1717 1339 1747">• invalid_protocol_msg</li> <li data-bbox="1058 1768 1307 1797">• invalid_session_id</li> <li data-bbox="1058 1818 1323 1848">• inactive_connection</li> <li data-bbox="1058 1869 1234 1898">• invalid_user</li> <li data-bbox="1058 1919 1282 1948">• server_removed</li> </ul>



Name	Type	Description
ip	string	IP address of the Vscan server.
node	<a href="#">node</a>	
state	string	Specifies the server connection state indicating if it is in the connected or disconnected state. The following is a list of the possible states: <ul style="list-style-type: none"> <li>• connected - Connected</li> <li>• disconnected - Disconnected</li> </ul> enum: <ul style="list-style-type: none"> <li>• connected</li> <li>• disconnected</li> </ul>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
type	string	Server type. The possible values are: <ul style="list-style-type: none"> <li>• primary - Primary server</li> <li>• backup - Backup server</li> </ul>
update_time	string	Specifies the time the server is in the connected or disconnected state.
vendor	string	Name of the connected virus-scanner vendor.
version	string	Version of the connected virus-scanner.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete a Vscan configuration

DELETE /protocols/vscan/{svm.uuid}

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.
- Any active Vscan On-Access policy must first be disabled 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

- [DOC /protocols/vscan](#)
- [DOC /protocols/vscan/{svm.uuid}/scanner-pools](#)

### Parameters

Name	Type	In	Required	Description
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
10027259 +	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 "sp1". Reason: "Failed to delete policy. Reason: policy must be disabled before being deleted.". Retry 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve the Vscan configuration for an SVM

```
GET /protocols/vscan/{svm.uuid}
```

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:

- There can be only one Vscan configuration enabled for an SVM at any time.

### 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

- [DOC /protocols/vscan](#)
- [DOC /protocols/vscan/{svm.uuid}/scanner-pools](#)

### 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
<a href="#">_links</a>	<a href="#">_links</a>	
cache_clear	boolean	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	Specifies whether or not Vscan is enabled on the SVM.
on_access_policies	array[ <a href="#">vscan_on_access</a> ]	
on_demand_policies	array[ <a href="#">vscan_on_demand</a> ]	
scanner_pools	array[ <a href="#">vscan_scanner_pool</a> ]	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example response

```
{
  "_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": [
        "/voll/cold-files/",
        "/voll/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": "primary",
    "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

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	<a href="#">href</a>	

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
only_execute_access	boolean	Scan only files opened with execute-access.
scan_readonly_volumes	boolean	Specifies whether or not read-only volume can be scanned.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

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.

Name	Type	Description
enabled	boolean	Status of the On-Access Vscan policy

Name	Type	Description
mandatory	boolean	Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning.
name	string	On-Access policy ame
scope	<a href="#">scope</a>	

schedule

Schedule of the task.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

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.

Name	Type	Description
log_path	string	The path from the Vserver root where the task report is created.
name	string	On-Demand task name
scan_paths	array[string]	List of paths that need to be scanned.
schedule	<a href="#">schedule</a>	Schedule of the task.
scope	<a href="#">scope</a>	

#### cluster\_reference

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

#### 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".

Name	Type	Description
cluster	<a href="#">cluster_reference</a>	
name	string	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 ".".

Name	Type	Description
privileged_users	array[string]	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. <ul style="list-style-type: none"> <li>example: ["cifs\u1", "cifs\u2"]</li> </ul>
role	string	Specifies the role of the scanner pool. The possible values are: <ul style="list-style-type: none"> <li>primary - Always active.</li> <li>secondary - Active only when none of the primary external virus-scanning servers are connected.</li> <li>idle - Always inactive.</li> </ul>
servers	array[string]	Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP. <ul style="list-style-type: none"> <li>example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]</li> </ul>

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the Vscan configuration for an SVM

PATCH /protocols/vscan/{svm.uuid}

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

- [DOC /protocols/vscan](#)
- [DOC /protocols/vscan/{svm.uuid}/scanner-pools](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

#### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
cache_clear	boolean	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	Specifies whether or not Vscan is enabled on the SVM.
on_access_policies	array[ <a href="#">vscan_on_access</a> ]	
on_demand_policies	array[ <a href="#">vscan_on_demand</a> ]	
scanner_pools	array[ <a href="#">vscan_scanner_pool</a> ]	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "_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": [
        "/voll/cold-files/",
        "/voll/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": "primary",
    "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"
}
}

```



## Response

Status: 200, Ok

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
10027015 +	Attempting to enable a Vscan but no active scanner-pool exists for the specified SVM
10027011 +	Attempting to enable a Vscan for an SVM for which no CIFS server exists
10027023 +	Attempting to enable a Vscan for an SVM for which no active Vscan On-Access policy exists

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	<a href="#">href</a>	

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
only_execute_access	boolean	Scan only files opened with execute-access.
scan_readonly_volumes	boolean	Specifies whether or not read-only volume can be scanned.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

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.

Name	Type	Description
enabled	boolean	Status of the On-Access Vscan policy

Name	Type	Description
mandatory	boolean	Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning.
name	string	On-Access policy ame
scope	<a href="#">scope</a>	

schedule

Schedule of the task.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

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.

Name	Type	Description
log_path	string	The path from the Vserver root where the task report is created.
name	string	On-Demand task name
scan_paths	array[string]	List of paths that need to be scanned.
schedule	<a href="#">schedule</a>	Schedule of the task.
scope	<a href="#">scope</a>	

#### cluster\_reference

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

#### 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".

Name	Type	Description
cluster	<a href="#">cluster_reference</a>	
name	string	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 ".".

Name	Type	Description
privileged_users	array[string]	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. <ul style="list-style-type: none"> <li>example: ["cifs\u1", "cifs\u2"]</li> </ul>
role	string	Specifies the role of the scanner pool. The possible values are: <ul style="list-style-type: none"> <li>primary - Always active.</li> <li>secondary - Active only when none of the primary external virus-scanning servers are connected.</li> <li>idle - Always inactive.</li> </ul>
servers	array[string]	Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP. <ul style="list-style-type: none"> <li>example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]</li> </ul>

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
cache_clear	boolean	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	Specifies whether or not Vscan is enabled on the SVM.
on_access_policies	array[ <a href="#">vscan_on_access</a> ]	
on_demand_policies	array[ <a href="#">vscan_on_demand</a> ]	
scanner_pools	array[ <a href="#">vscan_scanner_pool</a> ]	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage Vscan On-Access policies

### Protocols Vscan svm.uuid on-access-policies endpoint overview

#### Overview

Vscan On-Access scanning is used 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 CIFS share. The Vscan On-Access policy configuration defines the scope and status of On-Access scanning on file objects. This API is used 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"
```

---

## Retrieve a Vscan On-Access policy

GET /protocols/vscan/{svm.uuid}/on-access-policies

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

- [DOC /protocols/vscan/{svm.uuid}/on-access-policies](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
enabled	boolean	query	False	Filter by enabled
name	string	query	False	Filter by name
mandatory	boolean	query	False	Filter by mandatory
scope.include_extensions	string	query	False	Filter by scope.include_extensions
scope.scan_readonly_volumes	boolean	query	False	Filter by scope.scan_readonly_volumes
scope.exclude_extensions	string	query	False	Filter by scope.exclude_extensions
scope.max_file_size	integer	query	False	Filter by scope.max_file_size
scope.scan_without_extension	boolean	query	False	Filter by scope.scan_without_extension
scope.exclude_paths	string	query	False	Filter by scope.exclude_paths
scope.only_execute_access	boolean	query	False	Filter by scope.only_execute_access
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">vscan_on_access</a> ]	

## Example response

```
{
  "_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

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	<a href="#">href</a>	
self	<a href="#">href</a>	

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
only_execute_access	boolean	Scan only files opened with execute-access.
scan_readonly_volumes	boolean	Specifies whether or not read-only volume can be scanned.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

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.

Name	Type	Description
enabled	boolean	Status of the On-Access Vscan policy
mandatory	boolean	Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning.
name	string	On-Access policy ame
scope	<a href="#">scope</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create a Vscan On-Access policy

POST /protocols/vscan/{svm.uuid}/on-access-policies

Creates a Vscan On-Access policy. Created only on a data SVM.

**Important notes:**

- The policy needs to be enabled on an SVM before its files can be scanned.
- Only one On-Access policy can be enabled on an SVM at a time. 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

- [DOC /protocols/vscan/{svm.uuid}/on-access-policies](#)

## Parameters

Name	Type	In	Required	Description
<code>svm.uuid</code>	string	path	True	UUID of the SVM to which this object belongs.

## Request Body

Name	Type	Description
<code>enabled</code>	boolean	Status of the On-Access Vscan policy

Name	Type	Description
mandatory	boolean	Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning.
name	string	On-Access policy ame
scope	<a href="#">scope</a>	

### Example request

```
{
  "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

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">vscan_on_access</a> ]	

## Example response

```
{
  "_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

## 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

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
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

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

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
only_execute_access	boolean	Scan only files opened with execute-access.
scan_readonly_volumes	boolean	Specifies whether or not read-only volume can be scanned.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

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.

Name	Type	Description
enabled	boolean	Status of the On-Access Vscan policy
mandatory	boolean	Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning.
name	string	On-Access policy name
scope	<a href="#">scope</a>	

href



Name	Type	Description
href	string	

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete an antivirus On-Access policy configuration

DELETE /protocols/vscan/{svm.uuid}/on-access-policies/{name}

Deletes the anti-virus On-Access policy configuration.

### Related ONTAP commands

- `vserver vscan on-access-policy delete`

### Learn more

- [DOC /protocols/vscan/{svm.uuid}/on-access-policies](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	

### 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.

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve the Vscan On-Access policy configuration for an SVM

```
GET /protocols/vscan/{svm.uuid}/on-access-policies/{name}
```

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

- [DOC /protocols/vscan/{svm.uuid}/on-access-policies](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	
fields	array[string]	query	False	Specify the fields to return.

### Response

Status: 200, Ok

Name	Type	Description
enabled	boolean	Status of the On-Access Vscan policy
mandatory	boolean	Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning.
name	string	On-Access policy name
scope	scope	

## Example response

```
{
  "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

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

### scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
only_execute_access	boolean	Scan only files opened with execute-access.
scan_readonly_volumes	boolean	Specifies whether or not read-only volume can be scanned.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the Vscan On-Access policy configuration for an SVM

PATCH /protocols/vscan/{svm.uuid}/on-access-policies/{name}

Updates the Vscan On-Access policy configuration and/or enables/disables the Vscan On-Access policy of an SVM. Configurations for an On-Access policy associated with an administrative SVM cannot be modified, although the policy associated with an administrative SVM can be enabled or disabled.

### 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

- [DOC /protocols/vscan/{svm.uuid}/on-access-policies](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	

### Request Body

Name	Type	Description
enabled	boolean	Status of the On-Access Vscan policy
mandatory	boolean	Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning.
name	string	On-Access policy name



Name	Type	Description
scope	scope	

### Example request

```

{
  "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: 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.

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.
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.

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

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
only_execute_access	boolean	Scan only files opened with execute-access.
scan_readonly_volumes	boolean	Specifies whether or not read-only volume can be scanned.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

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.

Name	Type	Description
enabled	boolean	Status of the On-Access Vscan policy
mandatory	boolean	Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning.
name	string	On-Access policy name
scope	<a href="#">scope</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage Vscan On-Demand policies

### 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",
      "scope": {
        "max_file_size": 10737418240,
        "include_extensions": [
          "mp*"
        ],
        "scan_without_extension": true
      },
      "_links": {
        "self": {
          "href": "/api/protocols/vscan/86fbc414-f140-11e8-8e22-0050568e0945/on_demand_policies/policy2"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/protocols/vscan/86fbc414-f140-11e8-8e22-0050568e0945/on_demand_policies?fields=*"
    }
  }
}

```

```
# The API:
/api/protocols/vscan/{svm.uuid}/on_demand_policies/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-8e22-0050568e0945/on_demand_policies/on-demand-task" -H "accept: application/json"

# The response:
{
  "svm": {
    "uuid": "86fbc414-f140-11e8-8e22-0050568e0945",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/86fbc414-f140-11e8-8e22-0050568e0945"
      }
    }
  },
  "name": "on-demand-policy",
  "scan_paths": [
    "/voll/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\": \"/vol0/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\": [ \"/vol1/cifs/\" ], \"scope\": { \"include_extensions\": [ \"mp*\" ], \"scan_without_extension\": true }}"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "vs1"
      },
      "name": "on-demand-policy",
      "scan_paths": [
        "vol1/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"
```

---

## Retrieve a Vscan On-Demand policy

```
GET /protocols/vscan/{svm.uuid}/on-demand-policies
```

Retrieves the Vscan On-Demand policy.

### Related ONTAP commands

- `vserver vscan on-demand-task show`

### Learn more

- [DOC /protocols/vscan/{svm.uuid}/on-demand-policies](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
log_path	string	query	False	Filter by log_path
scope.include_extensions	string	query	False	Filter by scope.include_extensions
scope.exclude_extensions	string	query	False	Filter by scope.exclude_extensions
scope.max_file_size	integer	query	False	Filter by scope.max_file_size
scope.exclude_paths	string	query	False	Filter by scope.exclude_paths
scope.scan_without_extension	boolean	query	False	Filter by scope.scan_without_extension
schedule.uuid	string	query	False	Filter by schedule.uuid
schedule.name	string	query	False	Filter by schedule.name
name	string	query	False	Filter by name
scan_paths	string	query	False	Filter by scan_paths
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.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">vscan_on_demand</a> ]	

## Example response

```
{
  "_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

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

schedule

Schedule of the task.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.

Name	Type	Description
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

#### 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.

Name	Type	Description
log_path	string	The path from the Vserver root where the task report is created.
name	string	On-Demand task name
scan_paths	array[string]	List of paths that need to be scanned.
schedule	<a href="#">schedule</a>	Schedule of the task.
scope	<a href="#">scope</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create a Vscan On-Demand policy

POST /protocols/vscan/{svm.uuid}/on-demand-policies

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

- [DOC /protocols/vscan/{svm.uuid}/on-demand-policies](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

#### Request Body

Name	Type	Description
log_path	string	The path from the Vserver root where the task report is created.
name	string	On-Demand task name
scan_paths	array[string]	List of paths that need to be scanned.
schedule	<a href="#">schedule</a>	Schedule of the task.
scope	<a href="#">scope</a>	

## Example request

```
{
  "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
  }
}
```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records

Name	Type	Description
records	array[ <a href="#">vscan_on_demand</a> ]	

## Example response

```
{
  "_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

### 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.



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	<a href="#">href</a>	

schedule

Schedule of the task.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

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.

Name	Type	Description
log_path	string	The path from the Vserver root where the task report is created.
name	string	On-Demand task name
scan_paths	array[string]	List of paths that need to be scanned.
schedule	<a href="#">schedule</a>	Schedule of the task.
scope	<a href="#">scope</a>	

[\\_links](#)

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

[error\\_arguments](#)

Name	Type	Description
code	string	Argument code
message	string	Message argument

[error](#)

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete a Vscan On-Demand configuration

```
DELETE /protocols/vscan/{svm.uuid}/on-demand-policies/{name}
```

Deletes the Vscan On-Demand configuration.

## Related ONTAP commands

- `vserver vscan on-demand-task delete`

## Learn more

- [DOC /protocols/vscan/{svm.uuid}/on-demand-policies](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve the Vscan On-Demand configuration for an SVM

```
GET /protocols/vscan/{svm.uuid}/on-demand-policies/{name}
```

Retrieves the Vscan On-Demand configuration of an SVM.

### Related ONTAP commands

- `vserver vscan on-demand-task show`

### Learn more

- [DOC /protocols/vscan/{svm.uuid}/on-demand-policies](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	

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

### Response

Status: 200, Ok

Name	Type	Description
log_path	string	The path from the Vserver root where the task report is created.
name	string	On-Demand task name
scan_paths	array[string]	List of paths that need to be scanned.
schedule	<a href="#">schedule</a>	Schedule of the task.
scope	<a href="#">scope</a>	

## Example response

```
{
  "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

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	<a href="#">href</a>	

schedule

Schedule of the task.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Update the Vscan On-Demand policy configuration for an SVM

```
PATCH /protocols/vscan/{svm.uuid}/on-demand-policies/{name}
```

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

- [DOC /protocols/vscan/{svm.uuid}/on-demand-policies](#)

#### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	

## Request Body

Name	Type	Description
log_path	string	The path from the Vserver root where the task report is created.
name	string	On-Demand task name
scan_paths	array[string]	List of paths that need to be scanned.
schedule	<a href="#">schedule</a>	Schedule of the task.
scope	<a href="#">scope</a>	

## Example request

```
{
  "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
  }
}
```

## 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.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

schedule

Schedule of the task.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

scope

Name	Type	Description
exclude_extensions	array[string]	List of file extensions for which scanning is not performed.
exclude_paths	array[string]	List of file paths for which scanning must not be performed.
include_extensions	array[string]	List of file extensions to be scanned.
max_file_size	integer	Maximum file size, in bytes, allowed for scanning.
scan_without_extension	boolean	Specifies whether or not files without any extension can be scanned.

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.

Name	Type	Description
log_path	string	The path from the Vserver root where the task report is created.
name	string	On-Demand task name
scan_paths	array[string]	List of paths that need to be scanned.
schedule	<a href="#">schedule</a>	Schedule of the task.
scope	<a href="#">scope</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage Vscan scanner-pool configuration

### 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"
```

### Retrieve the Vscan scanner-pool configuration for an SVM

GET /protocols/vscan/{svm.uuid}/scanner-pools

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

- [DOC /protocols/vscan/{svm.uuid}/scanner-pools](#)

#### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Name	Type	In	Required	Description
privileged_users	string	query	False	Filter by privileged_users
name	string	query	False	Filter by name
cluster.uuid	string	query	False	Filter by cluster.uuid
cluster.name	string	query	False	Filter by cluster.name
servers	string	query	False	Filter by servers
role	string	query	False	Filter by role
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response



Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">vscan_scanner_pool</a> ]	

### Example response

```
{
  "_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": "primary",
    "servers": [
      "1.1.1.1",
      "10.72.204.27",
      "vmwin204-27.fsct.nb"
    ]
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

cluster\_reference

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

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".

Name	Type	Description
cluster	<a href="#">cluster_reference</a>	
name	string	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 ".".

Name	Type	Description
privileged_users	array[string]	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. <ul style="list-style-type: none"> <li>example: ["cifs\u1", "cifs\u2"]</li> </ul>
role	string	Specifies the role of the scanner pool. The possible values are: <ul style="list-style-type: none"> <li>primary - Always active.</li> <li>secondary - Active only when none of the primary external virus-scanning servers are connected.</li> <li>idle - Always inactive.</li> </ul>
servers	array[string]	Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP. <ul style="list-style-type: none"> <li>example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]</li> </ul>

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Create the Vscan scanner-pool configuration for an SVM

POST /protocols/vscan/{svm.uuid}/scanner-pools

Creates a Vscan scanner-pool configuration for a specified SVM. A scanner-pool can be created 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

- [DOC /protocols/vscan/{svm.uuid}/scanner-pools](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

## Request Body

Name	Type	Description
cluster	<a href="#">cluster_reference</a>	
name	string	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]	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. <ul style="list-style-type: none"><li>• example: ["cifs\u1", "cifs\u2"]</li></ul>

Name	Type	Description
role	string	Specifies the role of the scanner pool. The possible values are: <ul style="list-style-type: none"> <li>• primary - Always active.</li> <li>• secondary - Active only when none of the primary external virus-scanning servers are connected.</li> <li>• idle - Always inactive.</li> </ul>
servers	array[string]	Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP. <ul style="list-style-type: none"> <li>• example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]</li> </ul>

#### Example request

```
{
  "cluster": {
    "_links": {
      "self": {
        "href": "/api/resource/clusterlink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "scanner-1",
  "privileged_users": [
    "cifs\\u1",
    "cifs\\u2"
  ],
  "role": "primary",
  "servers": [
    "1.1.1.1",
    "10.72.204.27",
    "vmwin204-27.fsct.nb"
  ]
}
```

**Response**

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">vscan_scanner_pool</a> ]	



## Example response

```
{
  "_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": "primary",
    "servers": [
      "1.1.1.1",
      "10.72.204.27",
      "vmwin204-27.fsct.nb"
    ]
  }
}
```

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
10027086 +	The specified list of servers contain one or more entries that cannot be resolved
10027258 +	The specified cluster_name does not exist
10027256 +	The specified cluster_uuid does not exist
10027257 +	The specified cluster_name and cluster_uuid are valid but belong to different clusters
10027248 +	Scanner-pool created successfully but failed to activate
10027107 +	The list of privileged users or list of servers specified is empty
10027108 +	The list of privileged users specified contains an invalid entry
10027063 +	Attempting to modify a scanner-pool on an administrative SVM with a data 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

cluster\_reference

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

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".

Name	Type	Description
cluster	<a href="#">cluster_reference</a>	
name	string	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 ".".

Name	Type	Description
privileged_users	array[string]	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. <ul style="list-style-type: none"> <li>example: ["cifs\u1", "cifs\u2"]</li> </ul>
role	string	Specifies the role of the scanner pool. The possible values are: <ul style="list-style-type: none"> <li>primary - Always active.</li> <li>secondary - Active only when none of the primary external virus-scanning servers are connected.</li> <li>idle - Always inactive.</li> </ul>
servers	array[string]	Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP. <ul style="list-style-type: none"> <li>example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]</li> </ul>

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete a Vscan scanner-pool configuration

```
DELETE /protocols/vscan/{svm.uuid}/scanner-pools/{name}
```

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, Vscan must be disabled on the specified SVM so that all of the scanner-pools configured on that SVM can be deleted.

### Related ONTAP commands

- `vserver vscan scanner-pool delete`

### Learn more

- [DOC /protocols/vscan/{svm.uuid}/scanner-pools](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	

## Response

Status: 200, Ok

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
10027070 +	Attempting to delete a scanner-pool but it is the only active scanner-pool for a Vscan enabled on the SVM
10027064 +	Attempting to delete a scanner-pool with a data SVM which was created with an administrative SVM

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve the Vscan scanner-pool configuration for an SVM

```
GET /protocols/vscan/{svm.uuid}/scanner-pools/{name}
```

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

- [DOC /protocols/vscan/{svm.uuid}/scanner-pools](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

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

### Response

Status: 200, Ok

Name	Type	Description
cluster	<a href="#">cluster_reference</a>	
name	string	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]	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. <ul style="list-style-type: none"> <li>example: ["cifs\u1", "cifs\u2"]</li> </ul>
role	string	Specifies the role of the scanner pool. The possible values are: <ul style="list-style-type: none"> <li>primary - Always active.</li> <li>secondary - Active only when none of the primary external virus-scanning servers are connected.</li> <li>idle - Always inactive.</li> </ul>



Name	Type	Description
servers	array[string]	Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP. <ul style="list-style-type: none"> <li>example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]</li> </ul>

### Example response

```

{
  "cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "scanner-1",
  "privileged_users": [
    "cifs\\u1",
    "cifs\\u2"
  ],
  "role": "primary",
  "servers": [
    "1.1.1.1",
    "10.72.204.27",
    "vmwin204-27.fsct.nb"
  ]
}

```

### 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	<a href="#">href</a>	

cluster\_reference

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the Vscan scanner-pool configuration for an SVM

```
PATCH /protocols/vscan/{svm.uuid}/scanner-pools/{name}
```

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

- [DOC /protocols/vscan/{svm.uuid}/scanner-pools](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
name	string	path	True	

## Request Body

Name	Type	Description
cluster	<a href="#">cluster_reference</a>	
name	string	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 ".".

Name	Type	Description
privileged_users	array[string]	<p>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.</p> <ul style="list-style-type: none"> <li>• example: ["cifs\u1", "cifs\u2"]</li> </ul>
role	string	<p>Specifies the role of the scanner pool. The possible values are:</p> <ul style="list-style-type: none"> <li>• primary - Always active.</li> <li>• secondary - Active only when none of the primary external virus-scanning servers are connected.</li> <li>• idle - Always inactive.</li> </ul>
servers	array[string]	<p>Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.</p> <ul style="list-style-type: none"> <li>• example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]</li> </ul>

## Example request

```
{
  "cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "scanner-1",
  "privileged_users": [
    "cifs\\u1",
    "cifs\\u2"
  ],
  "role": "primary",
  "servers": [
    "1.1.1.1",
    "10.72.204.27",
    "vmwin204-27.fsct.nb"
  ]
}
```

## Response

Status: 200, Ok

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
10027258 +	The specified cluster_name does not exist
10027256 +	The specified cluster_uuid does not exist
10027257 +	The specified cluster_name and cluster_uuid are valid but belong to different clusters

Error Code	Description
10027248 +	Scanner-pool updated successfully but failed to apply the specified role
10027107 +	The list of privileged users or list of servers specified is empty
10027108 +	The list of privileged users specified contains an invalid entry
10027063 +	Attempting to modify a scanner-pool on an administrative SVM with a data 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

cluster\_reference

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

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".

Name	Type	Description
cluster	<a href="#">cluster_reference</a>	
name	string	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 ".".



Name	Type	Description
privileged_users	array[string]	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. <ul style="list-style-type: none"> <li>example: ["cifs\u1", "cifs\u2"]</li> </ul>
role	string	Specifies the role of the scanner pool. The possible values are: <ul style="list-style-type: none"> <li>primary - Always active.</li> <li>secondary - Active only when none of the primary external virus-scanning servers are connected.</li> <li>idle - Always inactive.</li> </ul>
servers	array[string]	Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP. <ul style="list-style-type: none"> <li>example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]</li> </ul>

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## NVMe

### 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**

## **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 [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"
    }
  }
}

```

## Retrieve NVMe interfaces

GET /protocols/nvme/interfaces

Retrieves NVMe interfaces.

## Related ONTAP commands

- `vserver nvme show-interface`

## Learn more

- [DOC /protocols/nvme/interfaces](#)

## Parameters

Name	Type	In	Required	Description
name	string	query	False	Filter by name
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
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
fc_interface.wwpn	string	query	False	Filter by fc_interface.wwpn
fc_interface.wwnn	string	query	False	Filter by fc_interface.wwnn
fc_interface.port.name	string	query	False	Filter by fc_interface.port.name
fc_interface.port.node.name	string	query	False	Filter by fc_interface.port.node.name
fc_interface.port.uuid	string	query	False	Filter by fc_interface.port.uuid
transport_address	string	query	False	Filter by transport_address
uuid	string	query	False	Filter by 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_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[nvme_interface]	



## Example response

```
{
  "_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-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, 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

node

The node on which the FC port is located.

Name	Type	Description
name	string	The name of the node on which the FC port is located.

port

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FC port.
node	<a href="#">node</a>	The node on which the FC port is located.
uuid	string	The unique identifier of the FC port.

fc\_interface

The attributes specific to a Fibre Channel-based NVMe interface.

Name	Type	Description
_links	<a href="#">_links</a>	
port	port	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	The WWNN (world wide node name) of the Fibre Channel NVMe interface.
wwpn	string	The WWPn (world wide port name) of the Fibre Channel NVMe interface.

#### node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### 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 [POST /network/fc/interfaces](#) . After creation, the interfaces are available via this interface.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
enabled	boolean	The administrative state of the NVMe interface.
fc_interface	<a href="#">fc_interface</a>	The attributes specific to a Fibre Channel-based NVMe interface.
name	string	The name of the NVMe interface.
node	<a href="#">node</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
transport_address	string	The transport address of the NVMe interface.
uuid	string	The unique identifier of the NVMe interface.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Retrieve an NVMe interface

GET /protocols/nvme/interfaces/{uuid}

Retrieves an NVMe interface.

## Related ONTAP commands

- `vserver nvme show-interface`

## Learn more

- [DOC /protocols/nvme/interfaces](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	The administrative state of the NVMe interface.
fc_interface	<a href="#">fc_interface</a>	The attributes specific to a Fibre Channel-based NVMe interface.
name	string	The name of the NVMe interface.
node	<a href="#">node</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
transport_address	string	The transport address of the NVMe interface.
uuid	string	The unique identifier of the NVMe interface.

## Example response

```
{
  "_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.

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	<a href="#">href</a>	

node

The node on which the FC port is located.

Name	Type	Description
name	string	The name of the node on which the FC port is located.

port

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FC port.
node	<a href="#">node</a>	The node on which the FC port is located.
uuid	string	The unique identifier of the FC port.

fc\_interface

The attributes specific to a Fibre Channel-based NVMe interface.

Name	Type	Description
_links	<a href="#">_links</a>	
port	<a href="#">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.

Name	Type	Description
wwnn	string	The WWNN (world wide node name) of the Fibre Channel NVMe interface.
wwpn	string	The WWPN (world wide port name) of the Fibre Channel NVMe interface.

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

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

## Manage NVMe services

### 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.

#### 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": "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"
        }
      }
    }
  ]
}
```

## 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/bfblbeb0-
dc69-11e8-b29f-005056bb7341' -H 'accept: application/hal+json'
```

## Retrieve NVMe services

GET /protocols/nvme/services

Retrieves NVMe services.

### Related ONTAP commands

- `vserver nvme show`

### Learn more

- [DOC /protocols/nvme/services](#)

### Parameters

Name	Type	In	Required	Description
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
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.



Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">nvme_service</a> ]	

## 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"
    }
  }
}
```

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
enabled	boolean	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 <i>true</i> (enabled) in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Create an NVMe service

POST /protocols/nvme/services

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

- [DOC /protocols/nvme/services](#)

## Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	<p>The administrative state of the NVMe service. The NVMe service can be disabled to block all NVMe connectivity to the SVM.</p> <p>This is optional in POST and PATCH. The default setting is <i>true</i> (enabled) in POST.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

## Response

```
Status: 201, Created
```

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[nvme_service]	

### 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"
    }
  }
}

```

### 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 <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.
72089650	An NVMe service already exists for the specified SVM.
5374893	The SVM is stopped. The SVM must be running to create an NVMe service.
72089900	An NVMe service cannot be creating in an SVM that is configured for a SAN protocol.

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	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 <i>true</i> (enabled) in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete an NVMe service

DELETE /protocols/nvme/services/{svm.uuid}

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

- [DOC /protocols/nvme/services](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	

## 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.

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an NVMe service

GET /protocols/nvme/services/{svm.uuid}

Retrieves an NVMe service.

### Related ONTAP commands

- `vserver nvme show`

### Learn more

- [DOC /protocols/nvme/services](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	The unique identifier of the SVM whose NVMe service is to be retrieved.
fields	array[string]	query	False	Specify the fields to return.

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	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 <i>true</i> (enabled) in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "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.

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update an NVMe service

PATCH /protocols/nvme/services/{svm.uuid}

Updates an NVMe service.

### Related ONTAP commands

- `vserver nvme modify`

### Learn more

- [DOC /protocols/nvme/services](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	The unique identifier of the SVM whose NVMe service is to be updated.

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	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 <i>true</i> (enabled) in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.



## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

## Response

Status: 200, Ok

## 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.
72089651	The supplied SVM does not have an NVMe service.
5374893	The SVM is stopped. The SVM must be running to create an NVMe service.

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	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 <i>true</i> (enabled) in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## View NVMe subsystem controllers

### 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/{subsystem.uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystem-controllers/14875240-2594-11e9-abde-00a098984313' -H 'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "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_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_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",
  "interface": {
    "name": "symmcon_lif_fcnvme_symmcon_fcnvme_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"
    }
  }
}

```

```
}  
}
```

## Retrieve NVMe subsystem controllers

GET /protocols/nvme/subsystem-controllers

Retrieves NVMe subsystem controllers.

### Related ONTAP commands

- `vserver nvme subsystem controller show`

### Learn more

- [DOC /protocols/nvme/subsystem-controllers](#)

### Parameters

Name	Type	In	Required	Description
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
host.transport_address	string	query	False	Filter by host.transport_address
host.nqn	string	query	False	Filter by host.nqn
host.id	string	query	False	Filter by host.id
subsystem.name	string	query	False	Filter by subsystem.name
subsystem.uuid	string	query	False	Filter by subsystem.uuid
admin_queue.depth	integer	query	False	Filter by admin_queue.depth
id	string	query	False	Filter by id

Name	Type	In	Required	Description
interface.name	string	query	False	Filter by interface.name
interface.transport_address	string	query	False	Filter by interface.transport_address
interface.uuid	string	query	False	Filter by interface.uuid
io_queue.count	integer	query	False	Filter by io_queue.count
io_queue.depth	integer	query	False	Filter by io_queue.depth
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">nvme_subsystem_controller</a> ]	

## Example response

```
{
  "_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": "falc5941-2593-11e9-94c4-00a0989a1c8e"
    },
    "io_queue": {
      "count": 0,
      "depth": {
      }
    },
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",

```

```
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "subsystem": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

admin\_queue

Name	Type	Description
depth	integer	The depth of the admin queue for the controller.

host

Properties of the connected host.

Name	Type	Description
id	string	The host identifier registered with the controller.
nqn	string	The NVMe qualified name of the host.
transport_address	string	The transport address of the host.

interface

The logical interface through which the host is connected.

Name	Type	Description
name	string	The name of the logical interface.

Name	Type	Description
transport_address	string	The transport address of the logical interface.
uuid	string	The unique identifier of the logical interface.

#### io\_queue

Properties of the I/O queues available to the controller.

Name	Type	Description
count	integer	The number of I/O queues available to the controller.
depth	array[integer]	The depths of the I/O queues.

#### node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

#### subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the NVMe subsystem.
uuid	string	The unique identifier of the NVMe subsystem.

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### 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.

Name	Type	Description
_links	<a href="#">_links</a>	
admin_queue	<a href="#">admin_queue</a>	
host	<a href="#">host</a>	Properties of the connected host.
id	string	The identifier of the subsystem controller. This field consists of 4 zero-filled hexadecimal digits followed by an 'h'.
interface	<a href="#">interface</a>	The logical interface through which the host is connected.
io_queue	<a href="#">io_queue</a>	Properties of the I/O queues available to the controller.
node	<a href="#">node</a>	
subsystem	<a href="#">subsystem</a>	An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Retrieve an NVMe subsystem controller

GET /protocols/nvme/subsystem-controllers/{system.uuid}/{id}

Retrieves an NVMe subsystem controller.

#### Related ONTAP commands

- `vserver nvme subsystem controller show`

#### Learn more

- [DOC /protocols/nvme/subsystem-controllers](#)

#### Parameters

Name	Type	In	Required	Description
system.uuid	string	path	True	The unique identifier of the NVMe subsystem.
id	string	path	True	The unique identifier of the NVMe subsystem controller.
fields	array[string]	query	False	Specify the fields to return.

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
admin_queue	<a href="#">admin_queue</a>	
host	<a href="#">host</a>	Properties of the connected host.
id	string	The identifier of the subsystem controller. This field consists of 4 zero-filled hexadecimal digits followed by an 'h'.
interface	<a href="#">interface</a>	The logical interface through which the host is connected.
io_queue	<a href="#">io_queue</a>	Properties of the I/O queues available to the controller.
node	<a href="#">node</a>	
subsystem	<a href="#">subsystem</a>	An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example response

```
{
  "_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": {
    }
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "subsystem": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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
72090001	The supplied subsystem identifier 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

admin\_queue

Name	Type	Description
depth	integer	The depth of the admin queue for the controller.

host

Properties of the connected host.

Name	Type	Description
id	string	The host identifier registered with the controller.
nqn	string	The NVMe qualified name of the host.
transport_address	string	The transport address of the host.

interface

The logical interface through which the host is connected.

Name	Type	Description
name	string	The name of the logical interface.
transport_address	string	The transport address of the logical interface.
uuid	string	The unique identifier of the logical interface.

io\_queue



Properties of the I/O queues available to the controller.

Name	Type	Description
count	integer	The number of I/O queues available to the controller.
depth	array[integer]	The depths of the I/O queues.

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the NVMe subsystem.
uuid	string	The unique identifier of the NVMe subsystem.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage NVMe subsystem maps

### 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/namespace1",
        "_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'

```

## Retrieve NVMe subsystem maps

GET /protocols/nvme/subsystem-maps

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 [DOC Requesting specific fields](#) to learn more.

- `anagrpId`

### Related ONTAP commands

- `vserver nvme subsystem map show`

### Learn more

- [DOC /protocols/nvme/subsystem-maps](#)

### Parameters

Name	Type	In	Required	Description
<code>svm.uuid</code>	string	query	False	Filter by <code>svm.uuid</code>
<code>svm.name</code>	string	query	False	Filter by <code>svm.name</code>
<code>nsid</code>	string	query	False	Filter by <code>nsid</code>
<code>anagrpId</code>	string	query	False	Filter by <code>anagrpId</code>
<code>subsystem.name</code>	string	query	False	Filter by <code>subsystem.name</code>
<code>subsystem.uuid</code>	string	query	False	Filter by <code>subsystem.uuid</code>
<code>namespace.nodeName</code>	string	query	False	Filter by <code>namespace.nodeName</code>
<code>namespace.nodeUUID</code>	string	query	False	Filter by <code>namespace.nodeUUID</code>
<code>namespace.name</code>	string	query	False	Filter by <code>namespace.name</code>

Name	Type	In	Required	Description
namespace.uuid	string	query	False	Filter by namespace.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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">nvme_subsystem_map</a> ]	

## Example response

```
{
  "_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"
        }
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "svm": {
```



```
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svml",
  "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

namespace

The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	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	<a href="#">node</a>	
uuid	string	The unique identifier of the NVMe namespace. Valid in POST.

subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of

## NVMe-connected hosts.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the NVMe subsystem.
uuid	string	The unique identifier of the NVMe subsystem.

## svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

## 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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
anagrpid	string	<p>The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.</p> <p>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
namespace	<a href="#">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	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	<a href="#">subsystem</a>	An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create an NVMe subsystem map

POST /protocols/nvme/subsystem-maps

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

- [DOC /protocols/nvme/subsystem-maps](#)

### Request Body

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>anagrpId</code>	string	<p>The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.</p> <p>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
<code>namespace</code>	<code>namespace</code>	<p>The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.</p>

Name	Type	Description
nsid	string	<p>The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.</p> <p>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".</p>
subsystem	subsystem	An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.
svm	svm	SVM, applies only to SVM-scoped objects.

**Example request**



```

{
  "_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"
      }
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}

```



## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">nvme_subsystem_map</a> ]	

## Example response

```
{
  "_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"
        }
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "svm": {
```

```

    "_links": {
      "self": {
        "href": "/api/resource/link"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}

```

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
72090006	The NVMe namespace specified by <code>namespace.uuid</code> was not found.
72090007	The NVMe namespace specified by <code>namespace.name</code> was not found.
72090005	The specified <code>namespace.uuid</code> and <code>namespace.name</code> refer to different NVMe namespaces.
72090001	The NVMe subsystem specified by <code>subsystem.uuid</code> was not found.
72090021	The NVMe subsystem specified by <code>subsystem.name</code> was not found.
72090020	The specified <code>subsystem.uuid</code> and <code>subsystem.name</code> refer to different NVMe subsystems.
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.

Error Code	Description
72089904	The aggregate must be given back to its home node prior to mapping the NVMe namespace it contains.

Name	Type	Description
error	<a href="#">error</a>	

### 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	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

namespace

The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	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	<a href="#">node</a>	
uuid	string	The unique identifier of the NVMe namespace. Valid in POST.

subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	The name of the NVMe subsystem.
uuid	string	The unique identifier of the NVMe subsystem.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
anagrpId	string	<p>The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.</p> <p>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
namespace	<a href="#">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	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	<a href="#">subsystem</a>	An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

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

### Delete an NVMe subsystem map

```
DELETE /protocols/nvme/subsystem-maps/{subsystem.uuid}/{namespace.uuid}
```

Deletes an NVMe subsystem map.

#### Related ONTAP commands

- `vserver nvme subsystem map delete`

#### Learn more

- [DOC /protocols/nvme/subsystem-maps](#)

#### Parameters

Name	Type	In	Required	Description
subsystem.uuid	string	path	True	
namespace.uuid	string	path	True	

#### Response

```
Status: 200, Ok
```

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">nvme_subsystem_map</a> ]	



## Example response

```
{
  "_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/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"
        }
      },
      "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.

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

namespace

The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	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	<a href="#">node</a>	
uuid	string	The unique identifier of the NVMe namespace. Valid in POST.

subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of

NVMe-connected hosts.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the NVMe subsystem.
uuid	string	The unique identifier of the NVMe subsystem.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
anagrpid	string	<p>The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.</p> <p>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
namespace	<a href="#">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	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	<a href="#">subsystem</a>	An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an NVMe subsystem map

GET /protocols/nvme/subsystem-maps/{subsystem.uuid}/{namespace.uuid}

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 [DOC Requesting specific fields](#) to learn more.

- `anagrpId`

### Related ONTAP commands

- `vserver nvme subsystem map show`

### Learn more

- [DOC /protocols/nvme/subsystem-maps](#)

### Parameters

Name	Type	In	Required	Description
subsystem.uuid	string	path	True	The unique identifier of the NVMe subsystem.
namespace.uuid	string	path	True	The unique identifier of the NVMe namespace.
fields	array[string]	query	False	Specify the fields to return.

### Response

```
Status: 200, Ok
```

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
anagrpId	string	<p>The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.</p> <p>The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
namespace	<a href="#">namespace</a>	<p>The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.</p>
nsid	string	<p>The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.</p> <p>The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".</p>
subsystem	<a href="#">subsystem</a>	<p>An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.</p>
svm	<a href="#">svm</a>	<p>SVM, applies only to SVM-scoped objects.</p>

**Example response**

A large, empty rectangular box with a thin, dashed border, occupying most of the page. It is intended for an example response.



```

{
  "_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"
      }
    },
    "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.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

namespace

The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	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	<a href="#">node</a>	
uuid	string	The unique identifier of the NVMe namespace. Valid in POST.

subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	The name of the NVMe subsystem.
uuid	string	The unique identifier of the NVMe subsystem.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage NVMe subsystems

### 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/hal+json' -d '{ "svm": { "name": "svm1" }, "name":
"subsystem1", "os_type": "linux" }'
```

---

### 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/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"
          }
        }
      },
      "uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
      "name": "subsystem1",
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystems/acde901a-a379-4a91-9ea6-
1b728ed6696f"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystems"
    }
  }
}

```

### 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/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"
          }
        }
      },
      "uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
      "name": "subsystem1",
      "os_type": "linux",
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystems/acde901a-a379-4a91-9ea6-
1b728ed6696f"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystems?os_type=linux"
    }
  }
}
```

---

## 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/hal+json'

# The response:
{
  "svm": {
    "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/a009a9e7-4081-b576-7575-ada21efcaf16"
      }
    }
  },
  "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
    }
  }
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystems/acde901a-a379-4a91-9ea6-
1b728ed6696f"
    }
  }
}

```

## 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/hal+json'
```

```
# The response:
```

```
{  
  "svm": {  
    "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",  
    "name": "svml",  
    "_links": {  
      "self": {  
        "href": "/api/svm/svms/a009a9e7-4081-b576-7575-ada21efcaf16"  
      }  
    }  
  },  
  "uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",  
  "name": "subsystem1",  
  "subsystem_maps": [  
    {  
      "anagrpid": "00000001h",  
      "namespace": {  
        "uuid": "eeaaca23-128d-4a7d-be4a-dc9106705799",  
        "name": "/vol/vol1/namespace1"  
      },  
      "_links": {  
        "self": {  
          "href": "/api/storage/namespaces/eeaaca23-128d-4a7d-be4a-  
dc9106705799"  
        }  
      }  
    },  
    {  
      "nsid": "00000001h"  
    },  
    {  
      "_links": {  
        "self": {  
          "href": "/api/protocols/nvme/subsystem_maps/acde901a-a379-4a91-  
9ea6-1b728ed6696f/eeaaca23-128d-4a7d-be4a-dc9106705799"  
        }  
      }  
    }  
  ],  
  {  
    "anagrpid": "00000002h",  
    "namespace": {  
      "uuid": "feaaca23-83a0-4a7d-beda-dc9106705799",  
      "name": "/vol/vol1/namespace2"  
    }  
  }  
}
```

```

    "_links": {
      "self": {
        "href": "/api/storage/namespaces/feaaca23-83a0-4a7d-beda-
dc9106705799"
      }
    },
    "nsid": "00000002h"
  },
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystem_maps/acde901a-a379-4a91-
9ea6-1b728ed6696f/feaaca23-83a0-4a7d-beda-dc9106705799"
    }
  }
}
]
"_links": {
  "self": {
    "href": "/api/protocols/nvme/subsystems/acde901a-a379-4a91-9ea6-
1b728ed6696f"
  }
}
}
}

```

### 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/hal+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/hal+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/hal+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/hal+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/hal+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/hal+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/hal+json'
```

```
# The response:
```

```
{
"records": [
  {
    "nqn": "nqn.1992-01.com.example:subsys1.host1",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystems/acde901a-a379-4a91-9ea6-
1b728ed6696f/hosts/nqn.1992-01.com.example%3Asubsys1.host1"
      }
    }
  },
```

```
{
  "nqn": "nqn.1992-01.com.example:subsys1.host2",
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystems/acde901a-a379-4a91-9ea6-1b728ed6696f/hosts/nqn.1992-01.com.example%3Asubsys1.host2"
    }
  }
},
{
  "nqn": "nqn.1992-01.com.example:subsys1.host3",
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystems/acde901a-a379-4a91-9ea6-1b728ed6696f/hosts/nqn.1992-01.com.example%3Asubsys1.host3"
    }
  }
}
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/protocols/nvme/subsystems/acde901a-a379-4a91-9ea6-1b728ed6696f/hosts"
  }
}
}
```

---

## 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/hal+json'

# The response:
{
  "subsystem": {
    "uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystems/acde901a-a379-4a91-9ea6-
1b728ed6696f"
      }
    }
  },
  "nqn": "nqn.1992-01.com.example:subsys1.host1",
  "io_queue": {
    "count": 4,
    "depth": 32
  },
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystems/acde901a-a379-4a91-9ea6-
1b728ed6696f/hosts/nqn.1992-01.com.example%3Asubsys1.host1"
    }
  }
}
```

---

## 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/hal+json'
```

## Retrieve NVMe subsystems

GET /protocols/nvme/subsystems

Retrieves NVMe subsystems.

### Related ONTAP commands

- `vserver nvme subsystem host show`
- `vserver nvme subsystem map show`
- `vserver nvme subsystem show`

### Learn more

- [DOC /protocols/nvme/subsystems](#)

### Parameters

Name	Type	In	Required	Description
serial_number	string	query	False	Filter by serial_number
io_queue.default.depth	integer	query	False	Filter by io_queue.default.depth
io_queue.default.count	integer	query	False	Filter by io_queue.default.count
os_type	string	query	False	Filter by os_type
comment	string	query	False	Filter by comment
target_nqn	string	query	False	Filter by target_nqn
uuid	string	query	False	Filter by uuid
subsystem_maps.ana_grpid	string	query	False	Filter by subsystem_maps.ana_grpid
subsystem_maps.namespace.name	string	query	False	Filter by subsystem_maps.namespace.name

Name	Type	In	Required	Description
subsystem_maps.namespace.uuid	string	query	False	Filter by subsystem_maps.namespace.uuid
subsystem_maps.nsid	string	query	False	Filter by subsystem_maps.nsid
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
name	string	query	False	Filter by name
hosts.nqn	string	query	False	Filter by hosts.nqn
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc



## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">nvme_subsystem</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "comment": "string",
    "hosts": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "nqn": "nqn.1992-01.example.com:string"
    },
    "io_queue": {
      "default": {
        "count": 4,
        "depth": 16
      }
    },
    "name": "subsystem1",
    "os_type": "hyper_v",
    "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

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

hosts

An NVMe host of an NVMe subsystem.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target.

default

The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

Name	Type	Description
count	integer	The number of host I/O queue pairs.
depth	integer	The host I/O queue depth.

io\_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

Name	Type	Description
default	<a href="#">default</a>	The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

#### namespace

An NVMe namespace mapped to the NVMe subsystem.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the NVMe namespace.
uuid	string	The unique identifier of the NVMe namespace.

#### subsystem\_maps

An NVMe namespace mapped to the NVMe subsystem.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
anagrpid	string	The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.  The format for an ANAGRPIP is 8 hexadecimal digits (zero-filled) followed by a lower case "h".
namespace	<a href="#">namespace</a>	An NVMe namespace mapped to the NVMe subsystem.
nsid	string	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

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### nvme\_subsystem

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

Name	Type	Description
_links	<a href="#">_links</a>	
comment	string	A configurable comment for the NVMe subsystem. Optional in POST and PATCH.
hosts	array[ <a href="#">hosts</a> ]	
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.
name	string	The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST.
os_type	string	The host operating system of the NVMe subsystem's hosts. Required in POST.
serial_number	string	The serial number of the NVMe subsystem.

Name	Type	Description
subsystem_maps	array[ <a href="#">subsystem_maps</a> ]	The NVMe namespaces mapped to the NVMe subsystem.  There is an added cost to retrieving property values for <code>subsystem_maps</code> . They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target_nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target.
uuid	string	The unique identifier of the NVMe subsystem.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.



## Create an NVMe subsystem

POST /protocols/nvme/subsystems

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

- [DOC /protocols/nvme/subsystems](#)

### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	A configurable comment for the NVMe subsystem. Optional in POST and PATCH.
hosts	array[ <a href="#">hosts</a> ]	
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.
name	string	The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST.
os_type	string	The host operating system of the NVMe subsystem's hosts. Required in POST.
serial_number	string	The serial number of the NVMe subsystem.

Name	Type	Description
subsystem_maps	array[ <a href="#">subsystem_maps</a> ]	<p>The NVMe namespaces mapped to the NVMe subsystem.</p> <p>There is an added cost to retrieving property values for <code>subsystem_maps</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target_nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target.
uuid	string	The unique identifier of the NVMe subsystem.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "hosts": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "nqn": "nqn.1992-01.example.com:string"
  },
  "io_queue": {
    "default": {
      "count": 4,
      "depth": 16
    }
  },
  "name": "subsystem1",
  "os_type": "hyper_v",
  "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"
}
```

## Response

Status: 201, Created

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">nvme_subsystem</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "comment": "string",
    "hosts": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "nqn": "nqn.1992-01.example.com:string"
    },
    "io_queue": {
      "default": {
        "count": 4,
        "depth": 16
      }
    },
    "name": "subsystem1",
    "os_type": "hyper_v",
    "serial_number": "wCVsgFMiuMhVAAAAAAB",
    "subsystem_maps": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "anagrpid": "00103050h",
      "namespace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      }
    }
  }
}
```

```

    }
  },
  "name": "/vol/vol1/namespacel",
  "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

## ONTAP Error Response Codes

Error Code	Description
2621462	The supplied SVM does not exist.
2621706	The specified <code>svm.uuid</code> and <code>svm.name</code> do not refer to the same SVM.
2621707	The <code>svm.uuid</code> or <code>svm.name</code> must be provided.
72090029	The NVMe service does not exist.
72090025	The NVMe subsystem already exists for the SVM.
72089709	The NVMe subsystem name contains an invalid character.

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	<a href="#">href</a>	

hosts

An NVMe host of an NVMe subsystem.

Name	Type	Description
_links	<a href="#">_links</a>	
nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target.

default

The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

Name	Type	Description
count	integer	The number of host I/O queue pairs.
depth	integer	The host I/O queue depth.

io\_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

Name	Type	Description
default	<a href="#">default</a>	The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

namespace

An NVMe namespace mapped to the NVMe subsystem.



Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the NVMe namespace.
uuid	string	The unique identifier of the NVMe namespace.

#### subsystem\_maps

An NVMe namespace mapped to the NVMe subsystem.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
anagrpId	string	The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.  The format for an ANAGRPI is 8 hexadecimal digits (zero-filled) followed by a lower case "h".
namespace	<a href="#">namespace</a>	An NVMe namespace mapped to the NVMe subsystem.
nsid	string	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

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### nvme\_subsystem

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

Name	Type	Description
_links	<a href="#">_links</a>	
comment	string	A configurable comment for the NVMe subsystem. Optional in POST and PATCH.
hosts	array[ <a href="#">hosts</a> ]	
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.
name	string	The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST.
os_type	string	The host operating system of the NVMe subsystem's hosts. Required in POST.
serial_number	string	The serial number of the NVMe subsystem.
subsystem_maps	array[ <a href="#">subsystem_maps</a> ]	<p>The NVMe namespaces mapped to the NVMe subsystem.</p> <p>There is an added cost to retrieving property values for <code>subsystem_maps</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target_nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target.

Name	Type	Description
uuid	string	The unique identifier of the NVMe subsystem.

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve NVMe subsystem hosts

GET /protocols/nvme/subsystems/{subsystem.uuid}/hosts

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 [DOC Requesting specific fields](#) to learn more.

- `subsystem_maps.*`

## Related ONTAP commands

- `vserver nvme subsystem map show`
- `vserver nvme subsystem show`

## Learn more

- [DOC /protocols/nvme/subsystems](#)

## Parameters

Name	Type	In	Required	Description
subsystem.uuid	string	path	True	The unique identifier of the NVMe subsystem.
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[nvme_subsystem_host]	

## Example response

```
{
  "_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",
    "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"
      }
    },
    "subsystem": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  }
}
```

```
    }
  },
  "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

io\_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

Name	Type	Description
count	integer	The number of I/O queue pairs. The default value is inherited from the owning NVMe subsystem.
depth	integer	The I/O queue depth. The default value is inherited from the owning NVMe subsystem.

subsystem

The NVMe subsystem to which the NVMe host has been provisioned.

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	The unique identifier of the NVMe subsystem.

records

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.



Name	Type	Description
_links	<a href="#">_links</a>	
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.
nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target. Not allowed in POST when the <code>records</code> property is used.
subsystem	<a href="#">subsystem</a>	The NVMe subsystem to which the NVMe host has been provisioned.

nvme\_subsystem\_host

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

Name	Type	Description
_links	<a href="#">_links</a>	
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.
nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target. Not allowed in POST when the <code>records</code> property is used.
records	array[ <a href="#">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	<a href="#">subsystem</a>	The NVMe subsystem to which the NVMe host has been provisioned.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Add NVMe subsystem hosts

POST `/protocols/nvme/subsystems/{subsystem.uuid}/hosts`

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

- [DOC /protocols/nvme/subsystems](#)

#### Parameters

Name	Type	In	Required	Description
subsystem.uuid	string	path	True	The unique identifier of the NVMe subsystem.

#### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.
nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target. Not allowed in POST when the <code>records</code> property is used.
records	array[ <a href="#">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	<a href="#">subsystem</a>	The NVMe subsystem to which the NVMe host has been provisioned.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "io_queue": {
    "count": 4,
    "depth": 32
  },
  "nqn": "nqn.1992-01.example.com:string",
  "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"
  }
},
"subsystem": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
}
```

## Response

Status: 201, Created

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">nvme_subsystem_host</a> ]	

## Example response

```
{
  "_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",
    "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"
      }
    },
    "subsystem": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  }
}
```

```

    }
  },
  "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 <i>nqn.1992-01.com.example:string</i> .
72089772	The NQN is invalid. Add the prefix 'nqn'. An example of a valid NQN is <i>nqn.1992-01.com.example:string</i> .
72089773	The NQN is invalid. The date field must be formatted <i>yyyy-mm</i> . An example of a valid NQN is <i>nqn.1992-01.com.example:string</i> .
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.
72090030	A partial success occurred while adding multiple NVMe subsystem hosts to an NVMe subsystem.

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	<a href="#">href</a>	

io\_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

Name	Type	Description
count	integer	The number of I/O queue pairs. The default value is inherited from the owning NVMe subsystem.
depth	integer	The I/O queue depth. The default value is inherited from the owning NVMe subsystem.

subsystem

The NVMe subsystem to which the NVMe host has been provisioned.

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	The unique identifier of the NVMe subsystem.

records

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.
nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target. Not allowed in POST when the <code>records</code> property is used.
subsystem	<a href="#">subsystem</a>	The NVMe subsystem to which the NVMe host has been provisioned.

nvme\_subsystem\_host

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

Name	Type	Description
_links	<a href="#">_links</a>	
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.
nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target. Not allowed in POST when the <code>records</code> property is used.
records	array[ <a href="#">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	<a href="#">subsystem</a>	The NVMe subsystem to which the NVMe host has been provisioned.

\_links

Name	Type	Description
next	<a href="#">href</a>	

Name	Type	Description
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete an NVMe subsystem host

```
DELETE /protocols/nvme/subsystems/{subsystem.uuid}/hosts/{nqn}
```

Deletes an NVMe subsystem host from an NVMe subsystem.

#### Related ONTAP commands

- `vserver nvme subsystem host remove`

#### Learn more

- [DOC /protocols/nvme/subsystems](#)

#### Parameters

Name	Type	In	Required	Description
subsystem.uuid	string	path	True	
nqn	string	path	True	

#### 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 <i>nqn.1992-01.com.example:string</i> .
72089772	The NQN is invalid. Add the prefix ' <i>nqn</i> '. An example of a valid NQN is <i>nqn.1992-01.com.example:string</i> .
72089773	The NQN is invalid. The date field must be formatted <i>yyyy-mm</i> . An example of a valid NQN is <i>nqn.1992-01.com.example:string</i> .

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an NVMe subsystem host

GET /protocols/nvme/subsystems/{subsystem.uuid}/hosts/{nqn}

Retrieves an NVMe subsystem host of an NVMe subsystem.

### Related ONTAP commands

- `vserver nvme subsystem host show`

### Learn more

- [DOC /protocols/nvme/subsystems](#)

### Parameters

Name	Type	In	Required	Description
subsystem.uuid	string	path	True	The unique identifier of the NVMe subsystem.
nqn	string	path	True	The NVMe qualified name (NQN) used to identify the NVMe subsystem host.

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

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.
nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target. Not allowed in POST when the <code>records</code> property is used.
records	array[ <a href="#">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	<a href="#">subsystem</a>	The NVMe subsystem to which the NVMe host has been provisioned.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "io_queue": {
    "count": 4,
    "depth": 32
  },
  "nqn": "nqn.1992-01.example.com:string",
  "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"
  }
},
"subsystem": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "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	<a href="#">href</a>	

io\_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

Name	Type	Description
count	integer	The number of I/O queue pairs. The default value is inherited from the owning NVMe subsystem.
depth	integer	The I/O queue depth. The default value is inherited from the owning NVMe subsystem.

subsystem

The NVMe subsystem to which the NVMe host has been provisioned.

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	The unique identifier of the NVMe subsystem.

records

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.
nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target. Not allowed in POST when the <code>records</code> property is used.
subsystem	<a href="#">subsystem</a>	The NVMe subsystem to which the NVMe host has been provisioned.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Remove an NVMe subsystem

```
DELETE /protocols/nvme/subsystems/{uuid}
```

Removes an NVMe subsystem.

### Related ONTAP commands

- `vserver nvme subsystem delete`

## Learn more

- [DOC /protocols/nvme/subsystems](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	
allow_delete_while_mapped	boolean	query	False	
allow_delete_with_hosts	boolean	query	False	

## 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 <code>allow_delete_while_mapped</code> query parameter to delete an NVMe subsystem with mapped NVMe namespaces.
72090024	The NVMe subsystem contains one or more NVMe hosts. Use the <code>allow_delete_with_hosts</code> query parameter to delete an NVMe subsystem with NVMe hosts.

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an NVMe subsystem

GET /protocols/nvme/subsystems/{uuid}

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 [DOC 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

- [DOC /protocols/nvme/subsystems](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier of the NVMe subsystem.
fields	array[string]	query	False	Specify the fields to return.

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	A configurable comment for the NVMe subsystem. Optional in POST and PATCH.
hosts	array[ <a href="#">hosts</a> ]	
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

Name	Type	Description
name	string	The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST.
os_type	string	The host operating system of the NVMe subsystem's hosts. Required in POST.
serial_number	string	The serial number of the NVMe subsystem.
subsystem_maps	array[ <a href="#">subsystem_maps</a> ]	<p>The NVMe namespaces mapped to the NVMe subsystem.</p> <p>There is an added cost to retrieving property values for <code>subsystem_maps</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target_nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target.
uuid	string	The unique identifier of the NVMe subsystem.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "hosts": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "nqn": "nqn.1992-01.example.com:string"
  },
  "io_queue": {
    "default": {
      "count": 4,
      "depth": 16
    }
  },
  "name": "subsystem1",
  "os_type": "hyper_v",
  "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

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	<a href="#">href</a>	

hosts

An NVMe host of an NVMe subsystem.

Name	Type	Description
_links	<a href="#">_links</a>	
nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target.

default

The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

Name	Type	Description
count	integer	The number of host I/O queue pairs.
depth	integer	The host I/O queue depth.

io\_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

Name	Type	Description
default	<a href="#">default</a>	The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

namespace

An NVMe namespace mapped to the NVMe subsystem.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the NVMe namespace.
uuid	string	The unique identifier of the NVMe namespace.

#### subsystem\_maps

An NVMe namespace mapped to the NVMe subsystem.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
anagrpid	string	The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.  The format for an ANAGRPIP is 8 hexadecimal digits (zero-filled) followed by a lower case "h".
namespace	<a href="#">namespace</a>	An NVMe namespace mapped to the NVMe subsystem.
nsid	string	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

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update an NVMe subsystem

PATCH /protocols/nvme/subsystems/{uuid}

Updates an NVMe subsystem.

### Related ONTAP commands

- `vserver nvme subsystem modify`

### Learn more

- [DOC /protocols/nvme/subsystems](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier of the NVMe subsystem.

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
comment	string	A configurable comment for the NVMe subsystem. Optional in POST and PATCH.
hosts	array[hosts]	
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.
name	string	The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST.
os_type	string	The host operating system of the NVMe subsystem's hosts. Required in POST.
serial_number	string	The serial number of the NVMe subsystem.
subsystem_maps	array[ <a href="#">subsystem_maps</a> ]	<p>The NVMe namespaces mapped to the NVMe subsystem.</p> <p>There is an added cost to retrieving property values for <code>subsystem_maps</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target_nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target.
uuid	string	The unique identifier of the NVMe subsystem.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "hosts": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "nqn": "nqn.1992-01.example.com:string"
  },
  "io_queue": {
    "default": {
      "count": 4,
      "depth": 16
    }
  },
  "name": "subsystem1",
  "os_type": "hyper_v",
  "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"
}
```

## 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

hosts

An NVMe host of an NVMe subsystem.

Name	Type	Description
_links	<a href="#">_links</a>	
nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target.

default

The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

Name	Type	Description
count	integer	The number of host I/O queue pairs.
depth	integer	The host I/O queue depth.

io\_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

Name	Type	Description
default	<a href="#">default</a>	The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

namespace

An NVMe namespace mapped to the NVMe subsystem.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the NVMe namespace.
uuid	string	The unique identifier of the NVMe namespace.

subsystem\_maps

An NVMe namespace mapped to the NVMe subsystem.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
anagrpId	string	The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.  The format for an ANAGRPI is 8 hexadecimal digits (zero-filled) followed by a lower case "h".
namespace	<a href="#">namespace</a>	An NVMe namespace mapped to the NVMe subsystem.
nsid	string	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

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

nvme\_subsystem



An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

Name	Type	Description
_links	<a href="#">_links</a>	
comment	string	A configurable comment for the NVMe subsystem. Optional in POST and PATCH.
hosts	array[ <a href="#">hosts</a> ]	
io_queue	<a href="#">io_queue</a>	The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.
name	string	The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST.
os_type	string	The host operating system of the NVMe subsystem's hosts. Required in POST.
serial_number	string	The serial number of the NVMe subsystem.
subsystem_maps	array[ <a href="#">subsystem_maps</a> ]	<p>The NVMe namespaces mapped to the NVMe subsystem.</p> <p>There is an added cost to retrieving property values for <code>subsystem_maps</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target_nqn	string	The NVMe qualified name (NQN) used to identify the NVMe storage target.

Name	Type	Description
uuid	string	The unique identifier of the NVMe subsystem.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage NVMe namespaces

### 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.

## Examples

### Creating an NVMe namespace

This example creates a 300 gigabyte NVMe namespace, with 4096-byte blocks, in SVM *svm1*, volume *vol1*, 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": "svm1" }, "os_type": "linux",
"space": { "block_size": "4096", "size": "300G" }, "name" :
"/vol/vol1/namespacel" }'

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "dccc3e6-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/namespacel",
      "location": {
        "namespace": "namespacel",
        "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=svml&status.state=online' -H 'accept:
application/hal+json'

# The response:
{
  "records": [
    {
      "uuid": "5c254d22-96a6-42ac-aad8-0cd9ebd126b6",
      "svm": {
        "name": "svml"
      },
      "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": "svml"
      },
      "name": "/vol/vol1/namespacel",
      "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",
    "anagrpid": "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"
  }
}
},
"_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 [/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 [/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'
```

## Retrieve NVMe namespaces

```
GET /storage/namespaces
```

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 [DOC Requesting specific fields](#) to learn more.

- `auto_delete`
- `subsystem_map.*`
- `status.mapped`

### Related ONTAP commands

- `vserver nvme namespace show`
- `vserver nvme subsystem map show`

### Learn more

- [DOC /storage/namespaces](#) to learn more and examples.

### Parameters

Name	Type	In	Required	Description
<code>status.container_state</code>	string	query	False	Filter by <code>status.container_state</code>
<code>status.state</code>	string	query	False	Filter by <code>status.state</code>
<code>status.read_only</code>	boolean	query	False	Filter by <code>status.read_only</code>
<code>status.mapped</code>	boolean	query	False	Filter by <code>status.mapped</code>
<code>svm.uuid</code>	string	query	False	Filter by <code>svm.uuid</code>
<code>svm.name</code>	string	query	False	Filter by <code>svm.name</code>
<code>name</code>	string	query	False	Filter by <code>name</code>
<code>space.size</code>	integer	query	False	Filter by <code>space.size</code>
<code>space.used</code>	integer	query	False	Filter by <code>space.used</code>

Name	Type	In	Required	Description
space.guarantee.reserved	boolean	query	False	Filter by space.guarantee.reserved
space.guarantee.requested	boolean	query	False	Filter by space.guarantee.requested
space.block_size	integer	query	False	Filter by space.block_size
subsystem_map.nsid	string	query	False	Filter by subsystem_map.nsid
subsystem_map.anagrp_id	string	query	False	Filter by subsystem_map.anagrp_id
subsystem_map.subsystem.name	string	query	False	Filter by subsystem_map.subsystem.name
subsystem_map.subsystem.uuid	string	query	False	Filter by subsystem_map.subsystem.uuid
enabled	boolean	query	False	Filter by enabled
auto_delete	boolean	query	False	Filter by auto_delete
os_type	string	query	False	Filter by os_type
location.namespace	string	query	False	Filter by location.namespace
location.volume.name	string	query	False	Filter by location.volume.name
location.volume.uuid	string	query	False	Filter by location.volume.uuid
location.qtree.id	integer	query	False	Filter by location.qtree.id

Name	Type	In	Required	Description
location.qtree.name	string	query	False	Filter by location.qtree.name
uuid	string	query	False	Filter by uuid
comment	string	query	False	Filter by comment
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.

Name	Type	Description
records	array[nvme_namespace]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "clone": {
      "source": {
        "name": "/vol/volume1/namespace1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "comment": "string",
    "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"
      }
    },
    "name": "/vol/volume1/qtree1/namespace1",
  }
}
```

```
"os_type": "hyper_v",
"space": {
  "block_size": 512,
  "size": 1073741824,
  "used": 0
},
"status": {
  "container_state": "online",
  "state": "online"
},
"subsystem_map": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "anagrpid": "00103050h",
  "nsid": "00000001h",
  "subsystem": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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.

Name	Type	Description
name	string	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	The unique identifier of the clone source NVMe namespace. Valid in POST and PATCH.

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`.

Name	Type	Description
<code>source</code>	<a href="#">source</a>	<p>The source NVMe namespace for a namespace clone operation. This can be specified using property <code>clone.source.uuid</code> or <code>clone.source.name</code>. If both properties are supplied, they must refer to the same namespace.</p> <p>Valid in POST to create a new NVMe namespace as a clone of the source.</p> <p>Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.</p>

#### 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.

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>id</code>	integer	The identifier for the qtree, unique within the qtree's volume.
<code>name</code>	string	The name of the qtree.

#### 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.

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

#### location

The location of the NVMe namespace within the ONTAP cluster. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

Name	Type	Description
namespace	string	<p>The base name component of the NVMe namespace. Valid in POST.</p> <p>If properties <code>name</code> and <code>location.namespace</code> are specified in the same request, they must refer to the base name.</p> <p>NVMe namespaces do not support rename.</p>
qtree	<a href="#">qtree</a>	<p>The qtree in which the NVMe namespace is optionally located. Valid in POST.</p> <p>If properties <code>name</code> and <code>location.qtree.name</code> and/or <code>location.qtree.uuid</code> are specified in the same request, they must refer to the same qtree.</p> <p>NVMe namespaces do not support rename.</p>

Name	Type	Description
volume	volume	<p>The volume in which the NVMe namespace is located. Valid in POST.</p> <p>If properties <code>name</code> and <code>location.volume.name</code> and/or <code>location.volume.uuid</code> are specified in the same request, they must refer to the same volume.</p> <p>NVMe namespaces do not support movement between volumes.</p>

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

Name	Type	Description
requested	boolean	<p>The requested space reservation policy for the NVMe namespace. If <i>true</i>, a space reservation is requested for the namespace; if <i>false</i>, 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%.</p> <p>The space reservation policy for an NVMe namespace is determined by ONTAP.</p> <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

Name	Type	Description
reserved	boolean	<p>Reports if the NVMe namespace is space guaranteed.</p> <p>This property is <i>true</i> if a space guarantee is requested and the containing volume and aggregate support the request. This property is <i>false</i> 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.</p>

space

The storage space related properties of the NVMe namespace.

Name	Type	Description
block_size	integer	<p>The size of blocks in the namespace in bytes.</p> <p>Valid in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.</p> <p>Valid in POST.</p>
guarantee	<a href="#">guarantee</a>	Properties that request and report the space guarantee for the NVMe namespace.
size	integer	<p>The total provisioned size of the NVMe namespace.</p> <p>NVMe namespaces do not support resize.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• example: 1073741824</li> <li>• readCreate: 1</li> </ul>

Name	Type	Description
used	integer	<p>The amount of space consumed by the main data stream of the NVMe namespace.</p> <p>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.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

status

Status information about the NVMe namespace.

Name	Type	Description
container_state	string	The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available.

Name	Type	Description
mapped	boolean	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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.
read_only	boolean	Reports if the NVMe namespace allows only read access.
state	string	The state of the NVMe namespace. Normal states for a namespace are <i>online</i> and <i>offline</i> . Other states indicate errors.

#### subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the NVMe subsystem.
uuid	string	The unique identifier of the NVMe subsystem.

#### 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 [DOC Requesting specific fields](#) to learn more.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
anagrpId	string	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	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	<a href="#">subsystem</a>	An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
auto_delete	boolean	<p>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.</p> <p>When set to <i>true</i>, 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.</p> <p>This property is optional in POST and PATCH. The default value for a new NVMe namespace is <i>false</i>.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
clone	clone	<p>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.</p> <p>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: <code>auto_delete</code> (unless specified in the request), <code>subsystem_map</code>, <code>status.state</code>, and <code>uuid</code>.</p>
comment	string	A configurable comment available for use by the administrator. Valid in POST and PATCH.
enabled	boolean	<p>The enabled state of the NVMe namespace. Namespaces can be disabled to prevent access to the namespace. Certain error conditions also cause the namespace to become disabled. If the namespace is disabled, you can consult the <code>state</code> property to determine if the namespace is administratively disabled (<i>offline</i>) or has become disabled as a result of an error. A namespace in an error condition can be brought online by setting the <code>enabled</code> property to <i>true</i> or brought administratively offline by setting the <code>enabled</code> property to <i>false</i>. Upon creation, an NVMe namespace is enabled by default. Valid in PATCH.</p>

Name	Type	Description
location	<a href="#">location</a>	<p>The location of the NVMe namespace within the ONTAP cluster. Valid in POST.</p> <p>NVMe namespaces do not support rename, or movement between volumes.</p> <ul style="list-style-type: none"> <li>• readCreate: 1</li> </ul>
name	string	<p>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.</p> <p>NVMe namespaces do not support rename, or movement between volumes.</p>
os_type	string	<p>The operating system type of the NVMe namespace.</p> <p>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.</p>
space	<a href="#">space</a>	<p>The storage space related properties of the NVMe namespace.</p>
status	<a href="#">status</a>	<p>Status information about the NVMe namespace.</p>

Name	Type	Description
subsystem_map	<a href="#">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 <code>subsystem_map</code> . They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the NVMe namespace.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create an NVMe namespace

POST `/storage/namespaces`

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

- [DOC /storage/namespaces](#)

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
auto_delete	boolean	<p>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.</p> <p>When set to <i>true</i>, 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.</p> <p>This property is optional in POST and PATCH. The default value for a new NVMe namespace is <i>false</i>.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
clone	<a href="#">clone</a>	<p>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.</p> <p>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: <code>auto_delete</code> (unless specified in the request), <code>subsystem_map</code>, <code>status.state</code>, and <code>uuid</code>.</p>

Name	Type	Description
comment	string	A configurable comment available for use by the administrator. Valid in POST and PATCH.
enabled	boolean	The enabled state of the NVMe namespace. Namespaces can be disabled to prevent access to the namespace. Certain error conditions also cause the namespace to become disabled. If the namespace is disabled, you can consult the <code>state</code> property to determine if the namespace is administratively disabled ( <i>offline</i> ) or has become disabled as a result of an error. A namespace in an error condition can be brought online by setting the <code>enabled</code> property to <i>true</i> or brought administratively offline by setting the <code>enabled</code> property to <i>false</i> . Upon creation, an NVMe namespace is enabled by default. Valid in PATCH.
location	location	<p>The location of the NVMe namespace within the ONTAP cluster. Valid in POST.</p> <p>NVMe namespaces do not support rename, or movement between volumes.</p> <ul style="list-style-type: none"> <li>• readCreate: 1</li> </ul>
name	string	<p>The fully qualified path name of the NVMe namespace composed of a <code>"/vol"</code> prefix, the volume name, the (optional) <code>qtree</code> name and base name of the namespace. Valid in POST.</p> <p>NVMe namespaces do not support rename, or movement between volumes.</p>

Name	Type	Description
os_type	string	<p>The operating system type of the NVMe namespace.</p> <p>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.</p>
space	<a href="#">space</a>	The storage space related properties of the NVMe namespace.
status	<a href="#">status</a>	Status information about the NVMe namespace.
subsystem_map	<a href="#">subsystem_map</a>	<p>The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.</p> <p>There is an added cost to retrieving property values for <code>subsystem_map</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the NVMe namespace.



## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "clone": {
    "source": {
      "name": "/vol/volume1/namespace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "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"
    }
  },
  "name": "/vol/volume1/qtree1/namespace1",
  "os_type": "hyper_v",
  "space": {
    "block_size": 512,
    "size": 1073741824,
    "used": 0
  },
  "status": {
    "container_state": "online",
    "state": "online"
  }
}
```

```

},
"subsystem_map": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"anagrpId": "00103050h",
"nsid": "00000001h",
"subsystem": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"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"
}

```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[nvme_namespace]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "clone": {
      "source": {
        "name": "/vol/volume1/namespace1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "comment": "string",
    "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"
      }
    },
    "name": "/vol/volume1/qtree1/namespace1",
  }
}
```

```
"os_type": "hyper_v",
"space": {
  "block_size": 512,
  "size": 1073741824,
  "used": 0
},
"status": {
  "container_state": "online",
  "state": "online"
},
"subsystem_map": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "anagrpid": "00103050h",
  "nsid": "00000001h",
  "subsystem": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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
2621462	The supplied 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.
72090009	An error occurred after successfully creating the NVMe namespace. Some properties were not set.
72090015	An error occurred after successfully creating the NVMe namespace preventing the retrieval of its properties.
72090005	The specified <code>clone.source.uuid</code> and <code>clone.source.name</code> do not refer to the same NVMe namespace.
72090006	The specified <code>clone.source</code> was not found.
72090007	The specified <code>clone.source</code> was not found.
917927	The specified volume was not found.
918236	The specified <code>location.volume.uuid</code> and <code>location.volume.name</code> do not refer to the same volume.
5374858	The volume specified by name is not the same as that specified by <code>location.volume</code> .
5242927	The specified <code>qtree</code> was not found.
5242950	The specified <code>location.qtree.id</code> and <code>location.qtree.name</code> do not refer to the same <code>qtree</code> .
5374860	The <code>qtree</code> specified by name is not the same as that specified by <code>location.qtree</code> .
5374861	The NVMe namespace base name specified by name is not the same as that specified by <code>location.name</code> .
13565952	The NVMe namespace clone request failed.
72090012	The property cannot be specified when creating an NVMe namespace clone. The <code>target</code> property of the error object identifies the property.
72090014	No volume was specified for the NVMe namespace.
72089721	The volume specified is in a load sharing mirror relationship. Namespaces are not supported in load sharing mirrors.

Error Code	Description
72090013	The property is required except when creating an NVMe namespace clone. The <code>target</code> property of the error object identifies the property.
5374862	No NVMe namespace path base name was provided for the namespace.
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.
5374352	An invalid name was provided 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 "}".
72089720	NVMe namespaces cannot be created in Snapshot copies.
72090033	The <code>clone.source.uuid</code> property is not supported when specifying a source NVMe namespace from a Snapshot copy.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

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.

Name	Type	Description
name	string	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	The unique identifier of the clone source NVMe namespace. Valid in POST and PATCH.

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`.



Name	Type	Description
source	source	<p>The source NVMe namespace for a namespace clone operation. This can be specified using property <code>clone.source.uuid</code> or <code>clone.source.name</code>. If both properties are supplied, they must refer to the same namespace.</p> <p>Valid in POST to create a new NVMe namespace as a clone of the source.</p> <p>Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.</p>

#### 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.

Name	Type	Description
_links	_links	
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree.

#### 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.

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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

location

The location of the NVMe namespace within the ONTAP cluster. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

Name	Type	Description
namespace	string	<p>The base name component of the NVMe namespace. Valid in POST.</p> <p>If <code>properties.name</code> and <code>location.namespace</code> are specified in the same request, they must refer to the base name.</p> <p>NVMe namespaces do not support rename.</p>
qtree	qtree	<p>The qtree in which the NVMe namespace is optionally located. Valid in POST.</p> <p>If <code>properties.name</code> and <code>location.qtree.name</code> and/or <code>location.qtree.uuid</code> are specified in the same request, they must refer to the same qtree.</p> <p>NVMe namespaces do not support rename.</p>

Name	Type	Description
volume	volume	<p>The volume in which the NVMe namespace is located. Valid in POST.</p> <p>If properties <code>name</code> and <code>location.volume.name</code> and/or <code>location.volume.uuid</code> are specified in the same request, they must refer to the same volume.</p> <p>NVMe namespaces do not support movement between volumes.</p>

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

Name	Type	Description
requested	boolean	<p>The requested space reservation policy for the NVMe namespace. If <i>true</i>, a space reservation is requested for the namespace; if <i>false</i>, 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%.</p> <p>The space reservation policy for an NVMe namespace is determined by ONTAP.</p> <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

Name	Type	Description
reserved	boolean	<p>Reports if the NVMe namespace is space guaranteed.</p> <p>This property is <i>true</i> if a space guarantee is requested and the containing volume and aggregate support the request. This property is <i>false</i> 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.</p>

space

The storage space related properties of the NVMe namespace.

Name	Type	Description
block_size	integer	<p>The size of blocks in the namespace in bytes.</p> <p>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.</p>
guarantee	<a href="#">guarantee</a>	Properties that request and report the space guarantee for the NVMe namespace.
size	integer	<p>The total provisioned size of the NVMe namespace.</p> <p>NVMe namespaces do not support resize.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• example: 1073741824</li> <li>• readCreate: 1</li> </ul>

Name	Type	Description
used	integer	<p>The amount of space consumed by the main data stream of the NVMe namespace.</p> <p>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.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

#### status

Status information about the NVMe namespace.

Name	Type	Description
container_state	string	The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available.

Name	Type	Description
mapped	boolean	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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.
read_only	boolean	Reports if the NVMe namespace allows only read access.
state	string	The state of the NVMe namespace. Normal states for a namespace are <i>online</i> and <i>offline</i> . Other states indicate errors.

#### subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
name	string	The name of the NVMe subsystem.
uuid	string	The unique identifier of the NVMe subsystem.

#### 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 [DOC Requesting specific fields](#) to learn more.

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
anagrpId	string	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	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	<a href="#">subsystem</a>	An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
auto_delete	boolean	<p>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.</p> <p>When set to <i>true</i>, 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.</p> <p>This property is optional in POST and PATCH. The default value for a new NVMe namespace is <i>false</i>.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>



Name	Type	Description
clone	clone	<p>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.</p> <p>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: <code>auto_delete</code> (unless specified in the request), <code>subsystem_map</code>, <code>status.state</code>, and <code>uuid</code>.</p>
comment	string	A configurable comment available for use by the administrator. Valid in POST and PATCH.
enabled	boolean	<p>The enabled state of the NVMe namespace. Namespaces can be disabled to prevent access to the namespace. Certain error conditions also cause the namespace to become disabled. If the namespace is disabled, you can consult the <code>state</code> property to determine if the namespace is administratively disabled (<i>offline</i>) or has become disabled as a result of an error. A namespace in an error condition can be brought online by setting the <code>enabled</code> property to <i>true</i> or brought administratively offline by setting the <code>enabled</code> property to <i>false</i>. Upon creation, an NVMe namespace is enabled by default. Valid in PATCH.</p>

Name	Type	Description
location	location	<p>The location of the NVMe namespace within the ONTAP cluster. Valid in POST.</p> <p>NVMe namespaces do not support rename, or movement between volumes.</p> <ul style="list-style-type: none"> <li>• readCreate: 1</li> </ul>
name	string	<p>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.</p> <p>NVMe namespaces do not support rename, or movement between volumes.</p>
os_type	string	<p>The operating system type of the NVMe namespace.</p> <p>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.</p>
space	space	<p>The storage space related properties of the NVMe namespace.</p>
status	status	<p>Status information about the NVMe namespace.</p>

Name	Type	Description
subsystem_map	<a href="#">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 <code>subsystem_map</code> . They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the NVMe namespace.

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

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

### Delete an NVMe namespace

```
DELETE /storage/namespaces/{uuid}
```

Deletes an NVMe namespace.

#### Related ONTAP commands

- `vserver nvme namespace delete`

#### Learn more

- [DOC /storage/namespaces](#)

#### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier of the NVMe namespace to delete.
allow_delete_while_mapped	boolean	query	False	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. <b>This parameter should be used with caution.</b> <ul style="list-style-type: none"> <li>• Default value:</li> </ul>

#### Response

Status: 200, Ok

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
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.
72090006	The specified namespace was not found.
72090007	The specified namespace 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an NVMe namespace

GET /storage/namespaces/{uuid}

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 [DOC Requesting specific fields](#) to learn more.

- `auto_delete`
- `subsystem_map.*`
- `status.mapped`

### Related ONTAP commands

- `vserver nvme namespace show`
- `vserver nvme subsystem map show`

### Learn more

- [DOC /storage/namespaces](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier of the NVMe namespace to retrieve.
fields	array[string]	query	False	Specify the fields to return.

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
auto_delete	boolean	<p>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.</p> <p>When set to <i>true</i>, 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.</p> <p>This property is optional in POST and PATCH. The default value for a new NVMe namespace is <i>false</i>.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
clone	clone	<p>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.</p> <p>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: <code>auto_delete</code> (unless specified in the request), <code>subsystem_map</code>, <code>status.state</code>, and <code>uuid</code>.</p>
comment	string	A configurable comment available for use by the administrator. Valid in POST and PATCH.
enabled	boolean	<p>The enabled state of the NVMe namespace. Namespaces can be disabled to prevent access to the namespace. Certain error conditions also cause the namespace to become disabled. If the namespace is disabled, you can consult the <code>state</code> property to determine if the namespace is administratively disabled (<i>offline</i>) or has become disabled as a result of an error. A namespace in an error condition can be brought online by setting the <code>enabled</code> property to <i>true</i> or brought administratively offline by setting the <code>enabled</code> property to <i>false</i>. Upon creation, an NVMe namespace is enabled by default. Valid in PATCH.</p>



Name	Type	Description
location	<a href="#">location</a>	<p>The location of the NVMe namespace within the ONTAP cluster. Valid in POST.</p> <p>NVMe namespaces do not support rename, or movement between volumes.</p> <ul style="list-style-type: none"> <li>• readCreate: 1</li> </ul>
name	string	<p>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.</p> <p>NVMe namespaces do not support rename, or movement between volumes.</p>
os_type	string	<p>The operating system type of the NVMe namespace.</p> <p>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.</p>
space	<a href="#">space</a>	<p>The storage space related properties of the NVMe namespace.</p>
status	<a href="#">status</a>	<p>Status information about the NVMe namespace.</p>

Name	Type	Description
subsystem_map	<a href="#">subsystem_map</a>	<p>The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.</p> <p>There is an added cost to retrieving property values for <code>subsystem_map</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the NVMe namespace.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "clone": {
    "source": {
      "name": "/vol/volume1/namespace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "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"
    }
  },
  "name": "/vol/volume1/qtree1/namespace1",
  "os_type": "hyper_v",
  "space": {
    "block_size": 512,
    "size": 1073741824,
    "used": 0
  },
  "status": {
    "container_state": "online",
    "state": "online"
  }
}
```

```

},
"subsystem_map": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"anagrpid": "00103050h",
"nsid": "00000001h",
"subsystem": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"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.

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	<a href="#">href</a>	

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.

Name	Type	Description
name	string	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	The unique identifier of the clone source NVMe namespace. Valid in POST and PATCH.

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`.

Name	Type	Description
source	source	<p>The source NVMe namespace for a namespace clone operation. This can be specified using property <code>clone.source.uuid</code> or <code>clone.source.name</code>. If both properties are supplied, they must refer to the same namespace.</p> <p>Valid in POST to create a new NVMe namespace as a clone of the source.</p> <p>Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.</p>

### 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.

Name	Type	Description
_links	_links	
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree.

### 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.

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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

location

The location of the NVMe namespace within the ONTAP cluster. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

Name	Type	Description
namespace	string	<p>The base name component of the NVMe namespace. Valid in POST.</p> <p>If properties <code>name</code> and <code>location.namespace</code> are specified in the same request, they must refer to the base name.</p> <p>NVMe namespaces do not support rename.</p>
qtree	qtree	<p>The qtree in which the NVMe namespace is optionally located. Valid in POST.</p> <p>If properties <code>name</code> and <code>location.qtree.name</code> and/or <code>location.qtree.uuid</code> are specified in the same request, they must refer to the same qtree.</p> <p>NVMe namespaces do not support rename.</p>



Name	Type	Description
volume	volume	<p>The volume in which the NVMe namespace is located. Valid in POST.</p> <p>If properties <code>name</code> and <code>location.volume.name</code> and/or <code>location.volume.uuid</code> are specified in the same request, they must refer to the same volume.</p> <p>NVMe namespaces do not support movement between volumes.</p>

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

Name	Type	Description
requested	boolean	<p>The requested space reservation policy for the NVMe namespace. If <i>true</i>, a space reservation is requested for the namespace; if <i>false</i>, 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%.</p> <p>The space reservation policy for an NVMe namespace is determined by ONTAP.</p> <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

Name	Type	Description
reserved	boolean	<p>Reports if the NVMe namespace is space guaranteed.</p> <p>This property is <i>true</i> if a space guarantee is requested and the containing volume and aggregate support the request. This property is <i>false</i> 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.</p>

space

The storage space related properties of the NVMe namespace.

Name	Type	Description
block_size	integer	<p>The size of blocks in the namespace in bytes.</p> <p>Valid in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.</p> <p>Valid in POST.</p>
guarantee	<a href="#">guarantee</a>	<p>Properties that request and report the space guarantee for the NVMe namespace.</p>
size	integer	<p>The total provisioned size of the NVMe namespace.</p> <p>NVMe namespaces do not support resize.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• example: 1073741824</li> <li>• readCreate: 1</li> </ul>

Name	Type	Description
used	integer	<p>The amount of space consumed by the main data stream of the NVMe namespace.</p> <p>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.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

#### status

Status information about the NVMe namespace.

Name	Type	Description
container_state	string	The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available.

Name	Type	Description
mapped	boolean	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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.
read_only	boolean	Reports if the NVMe namespace allows only read access.
state	string	The state of the NVMe namespace. Normal states for a namespace are <i>online</i> and <i>offline</i> . Other states indicate errors.

#### subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
name	string	The name of the NVMe subsystem.
uuid	string	The unique identifier of the NVMe subsystem.

#### 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 [DOC Requesting specific fields](#) to learn more.

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
anagrpid	string	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	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	<a href="#">subsystem</a>	An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments

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

### Update an NVMe namespace

```
PATCH /storage/namespaces/{uuid}
```

Updates an NVMe namespace.

#### Related ONTAP commands

- `volume file clone autodelete`
- `vserver nvme namespace modify`

#### Learn more

- [DOC /storage/namespaces](#)

#### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier of the NVMe namespace to update.

#### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
auto_delete	boolean	<p>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.</p> <p>When set to <i>true</i>, 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.</p> <p>This property is optional in POST and PATCH. The default value for a new NVMe namespace is <i>false</i>.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
clone	<a href="#">clone</a>	<p>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.</p> <p>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: <code>auto_delete</code> (unless specified in the request), <code>subsystem_map</code>, <code>status.state</code>, and <code>uuid</code>.</p>

Name	Type	Description
comment	string	A configurable comment available for use by the administrator. Valid in POST and PATCH.
enabled	boolean	The enabled state of the NVMe namespace. Namespaces can be disabled to prevent access to the namespace. Certain error conditions also cause the namespace to become disabled. If the namespace is disabled, you can consult the <code>state</code> property to determine if the namespace is administratively disabled ( <i>offline</i> ) or has become disabled as a result of an error. A namespace in an error condition can be brought online by setting the <code>enabled</code> property to <i>true</i> or brought administratively offline by setting the <code>enabled</code> property to <i>false</i> . Upon creation, an NVMe namespace is enabled by default. Valid in PATCH.
location	location	<p>The location of the NVMe namespace within the ONTAP cluster. Valid in POST.</p> <p>NVMe namespaces do not support rename, or movement between volumes.</p> <ul style="list-style-type: none"> <li>• readCreate: 1</li> </ul>
name	string	<p>The fully qualified path name of the NVMe namespace composed of a <code>"/vol"</code> prefix, the volume name, the (optional) <code>qtree</code> name and base name of the namespace. Valid in POST.</p> <p>NVMe namespaces do not support rename, or movement between volumes.</p>



Name	Type	Description
os_type	string	<p>The operating system type of the NVMe namespace.</p> <p>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.</p>
space	<a href="#">space</a>	The storage space related properties of the NVMe namespace.
status	<a href="#">status</a>	Status information about the NVMe namespace.
subsystem_map	<a href="#">subsystem_map</a>	<p>The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.</p> <p>There is an added cost to retrieving property values for <code>subsystem_map</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the NVMe namespace.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "clone": {
    "source": {
      "name": "/vol/volume1/namespace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "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"
    }
  },
  "name": "/vol/volume1/qtree1/namespace1",
  "os_type": "hyper_v",
  "space": {
    "block_size": 512,
    "size": 1073741824,
    "used": 0
  },
  "status": {
    "container_state": "online",
    "state": "online"
  }
}
```

```
},
"subsystem_map": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "anagrpid": "00103050h",
  "nsid": "00000001h",
  "subsystem": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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"
}
```

#### Response

Status: 200, Ok

#### Error

Status: Default

#### ONTAP Error Response Codes

Error Code	Description
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.
72090005	The specified <code>clone.source.uuid</code> and <code>clone.source.name</code> do not refer to the same LUN.
72090006	The specified namespace was not found. This can apply to <code>clone.source</code> or the target namespace. The <code>target</code> property of the error object identifies the property.
72090007	The specified namespace was not found. This can apply to <code>clone.source</code> or the target namespace. The <code>target</code> property of the error object identifies the property.
13565952	The namespace clone request failed.

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	<a href="#">href</a>	

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.

Name	Type	Description
name	string	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	The unique identifier of the clone source NVMe namespace. Valid in POST and PATCH.

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`.

Name	Type	Description
source	source	<p>The source NVMe namespace for a namespace clone operation. This can be specified using property <code>clone.source.uuid</code> or <code>clone.source.name</code>. If both properties are supplied, they must refer to the same namespace.</p> <p>Valid in POST to create a new NVMe namespace as a clone of the source.</p> <p>Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.</p>

#### 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.

Name	Type	Description
_links	_links	
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree.

#### 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.

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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

location

The location of the NVMe namespace within the ONTAP cluster. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

Name	Type	Description
namespace	string	<p>The base name component of the NVMe namespace. Valid in POST.</p> <p>If properties <code>name</code> and <code>location.namespace</code> are specified in the same request, they must refer to the base name.</p> <p>NVMe namespaces do not support rename.</p>
qtree	qtree	<p>The qtree in which the NVMe namespace is optionally located. Valid in POST.</p> <p>If properties <code>name</code> and <code>location.qtree.name</code> and/or <code>location.qtree.uuid</code> are specified in the same request, they must refer to the same qtree.</p> <p>NVMe namespaces do not support rename.</p>

Name	Type	Description
volume	volume	<p>The volume in which the NVMe namespace is located. Valid in POST.</p> <p>If properties <code>name</code> and <code>location.volume.name</code> and/or <code>location.volume.uuid</code> are specified in the same request, they must refer to the same volume.</p> <p>NVMe namespaces do not support movement between volumes.</p>

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

Name	Type	Description
requested	boolean	<p>The requested space reservation policy for the NVMe namespace. If <i>true</i>, a space reservation is requested for the namespace; if <i>false</i>, 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%.</p> <p>The space reservation policy for an NVMe namespace is determined by ONTAP.</p> <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>



Name	Type	Description
reserved	boolean	<p>Reports if the NVMe namespace is space guaranteed.</p> <p>This property is <i>true</i> if a space guarantee is requested and the containing volume and aggregate support the request. This property is <i>false</i> 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.</p>

## space

The storage space related properties of the NVMe namespace.

Name	Type	Description
block_size	integer	<p>The size of blocks in the namespace in bytes.</p> <p>Valid in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.</p> <p>Valid in POST.</p>
guarantee	<a href="#">guarantee</a>	Properties that request and report the space guarantee for the NVMe namespace.
size	integer	<p>The total provisioned size of the NVMe namespace.</p> <p>NVMe namespaces do not support resize.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• example: 1073741824</li> <li>• readCreate: 1</li> </ul>

Name	Type	Description
used	integer	<p>The amount of space consumed by the main data stream of the NVMe namespace.</p> <p>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.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

#### status

Status information about the NVMe namespace.

Name	Type	Description
container_state	string	The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available.

Name	Type	Description
mapped	boolean	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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.
read_only	boolean	Reports if the NVMe namespace allows only read access.
state	string	The state of the NVMe namespace. Normal states for a namespace are <i>online</i> and <i>offline</i> . Other states indicate errors.

#### subsystem

An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the NVMe subsystem.
uuid	string	The unique identifier of the NVMe subsystem.

#### 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 [DOC Requesting specific fields](#) to learn more.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
anagrpid	string	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	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	<a href="#">subsystem</a>	An NVMe subsystem maintains configuration state and NVMe namespace access control for a set of NVMe-connected hosts.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
auto_delete	boolean	<p>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.</p> <p>When set to <i>true</i>, 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.</p> <p>This property is optional in POST and PATCH. The default value for a new NVMe namespace is <i>false</i>.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
clone	clone	<p>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.</p> <p>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: <code>auto_delete</code> (unless specified in the request), <code>subsystem_map</code>, <code>status.state</code>, and <code>uuid</code>.</p>
comment	string	A configurable comment available for use by the administrator. Valid in POST and PATCH.
enabled	boolean	<p>The enabled state of the NVMe namespace. Namespaces can be disabled to prevent access to the namespace. Certain error conditions also cause the namespace to become disabled. If the namespace is disabled, you can consult the <code>state</code> property to determine if the namespace is administratively disabled (<i>offline</i>) or has become disabled as a result of an error. A namespace in an error condition can be brought online by setting the <code>enabled</code> property to <i>true</i> or brought administratively offline by setting the <code>enabled</code> property to <i>false</i>. Upon creation, an NVMe namespace is enabled by default. Valid in PATCH.</p>

Name	Type	Description
location	<a href="#">location</a>	<p>The location of the NVMe namespace within the ONTAP cluster. Valid in POST.</p> <p>NVMe namespaces do not support rename, or movement between volumes.</p> <ul style="list-style-type: none"> <li>• readCreate: 1</li> </ul>
name	string	<p>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.</p> <p>NVMe namespaces do not support rename, or movement between volumes.</p>
os_type	string	<p>The operating system type of the NVMe namespace.</p> <p>Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.</p>
space	<a href="#">space</a>	<p>The storage space related properties of the NVMe namespace.</p>
status	<a href="#">status</a>	<p>Status information about the NVMe namespace.</p>

Name	Type	Description
subsystem_map	<a href="#">subsystem_map</a>	<p>The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.</p> <p>There is an added cost to retrieving property values for <code>subsystem_map</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the NVMe namespace.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.



# Name-services

## 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

### 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 [/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"
```

## Retrieve DNS configurations for all SVMs

GET /name-services/dns

Retrieves the DNS configurations of all data SVMs.  
DNS configuration for the cluster is retrieved and managed via [/api/cluster](#) .

### Related ONTAP commands

- `vserver services name-service dns show`
- `vserver services name-service dns check`

### Learn more

- [DOC /name-services/dns](#)

### Parameters

Name	Type	In	Required	Description
servers	string	query	False	Filter by servers
domains	string	query	False	Filter by domains
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.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of DNS domain records
records	array[ <a href="#">dns</a> ]	

## Example response

```
{
  "_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

Name	Type	Description
error	<a href="#">error</a>	



## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

dns

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>
servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
svm	<a href="#">svm</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create DNS domain and server configurations

POST `/name-services/dns`

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:

1. The server is not a DNS server.
2. The server does not exist.
3. The server is unreachable.

### Learn more

- [DOC /name-services/dns](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>
servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
svm	<a href="#">svm</a>	

## Example request

```
{
  "_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"
  }
}
```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of DNS domain records
records	array[ <a href="#">dns</a> ]	

## Example response

```
{
  "_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

## 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
23724130	Cannot use an IPv6 name server address because there are no IPv6 LIFs

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	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

dns

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>
servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
svm	<a href="#">svm</a>	

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete a DNS domain configuration

```
DELETE /name-services/dns/{svm.uuid}
```

Deletes DNS domain configuration of the specified SVM.

#### Related ONTAP commands

- `vserver services name-service dns delete`

#### Learn more

- [DOC /name-services/dns](#)

#### Parameters

Name	Type	In	Required	Description
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	<a href="#">error</a>	

### 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve DNS domain and server configurations

GET /name-services/dns/{svm.uuid}

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 [/api/cluster](#) .

### Related ONTAP commands

- `vserver services name-service dns show`
- `vserver services name-service dns check`

### Learn more

- [DOC /name-services/dns](#)

### 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
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>
servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
svm	<a href="#">svm</a>	

## Example response

```
{
  "_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

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	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update DNS domain and server configurations

PATCH /name-services/dns/{svm.uuid}

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:

1. The server is not a DNS server.
2. The server does not exist.
3. The server is unreachable.

#### Learn more

- [DOC /name-services/dns](#)

#### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>
servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
svm	<a href="#">svm</a>	

## Example request

```
{
  "_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"
  }
}
```

## 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

Error Code	Description
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

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	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

dns

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>
servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
svm	<a href="#">svm</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error



Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage LDAP server configurations

### 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"
```

## Retrieve an LDAP configuration for all SVMs

GET /name-services/ldap

Retrieves the LDAP configurations for all SVMs.

### Learn more

- [DOC /name-services/ldap](#)

### 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	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok



Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of LDAP records.
records	array[ <a href="#">ldap_service</a> ]	

### Example response

```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "base_scope": "base",
    "min_bind_level": "anonymous",
    "port": 389,
    "preferred_ad_servers": {
    },
    "servers": {
    },
    "session_security": "none",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ldap\_service

Name	Type	Description
_links	<a href="#">_links</a>	
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.

Name	Type	Description
base_scope	string	Specifies the default search scope for LDAP queries: <ul style="list-style-type: none"> <li>• base - search the named entry only</li> <li>• onelevel - search all entries immediately below the DN</li> <li>• subtree - search the named DN entry and the entire subtree below the DN</li> </ul>
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"> <li>• anonymous - anonymous bind</li> <li>• simple - simple bind</li> <li>• sasl - Simple Authentication and Security Layer (SASL) bind</li> </ul>
port	integer	The port used to connect to the LDAP Servers.
preferred_ad_servers	array[string]	
schema	string	The name of the schema template used by the SVM. <ul style="list-style-type: none"> <li>• AD-IDMU - Active Directory Identity Management for UNIX</li> <li>• AD-SFU - Active Directory Services for UNIX</li> <li>• MS-AD-BIS - Active Directory Identity Management for UNIX</li> <li>• RFC-2307 - Schema based on RFC 2307</li> <li>• Custom schema</li> </ul>

Name	Type	Description
servers	array[string]	
session_security	string	Specifies the level of security to be used for LDAP communications: <ul style="list-style-type: none"> <li>• none - no signing or sealing</li> <li>• sign - sign LDAP traffic</li> <li>• seal - seal and sign LDAP traffic</li> </ul>
svm	<a href="#">svm</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Create an LDAP configuration for an SVM

POST /name-services/ldap

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.
- 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:

1. The server does not have LDAP installed.
2. The server or Active Directory domain is invalid.
3. The server or Active Directory domain is unreachable.

## Learn more

- [DOC /name-services/ldap](#)

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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.

Name	Type	Description
base_scope	string	Specifies the default search scope for LDAP queries: <ul style="list-style-type: none"> <li>• base - search the named entry only</li> <li>• onelevel - search all entries immediately below the DN</li> <li>• subtree - search the named DN entry and the entire subtree below the DN</li> </ul>
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"> <li>• anonymous - anonymous bind</li> <li>• simple - simple bind</li> <li>• sasl - Simple Authentication and Security Layer (SASL) bind</li> </ul>
port	integer	The port used to connect to the LDAP Servers.
preferred_ad_servers	array[string]	
schema	string	The name of the schema template used by the SVM. <ul style="list-style-type: none"> <li>• AD-IDMU - Active Directory Identity Management for UNIX</li> <li>• AD-SFU - Active Directory Services for UNIX</li> <li>• MS-AD-BIS - Active Directory Identity Management for UNIX</li> <li>• RFC-2307 - Schema based on RFC 2307</li> <li>• Custom schema</li> </ul>
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"> <li>• none - no signing or sealing</li> <li>• sign - sign LDAP traffic</li> <li>• seal - seal and sign LDAP traffic</li> </ul>
svm	svm	
use_start_tls	boolean	Specifies whether or not to use Start TLS over LDAP connections.

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "base_scope": "base",
  "min_bind_level": "anonymous",
  "port": 389,
  "preferred_ad_servers": {
  },
  "servers": {
  },
  "session_security": "none",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
```

### Response



Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of LDAP records.
records	array[ <a href="#">ldap_service</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "base_scope": "base",
    "min_bind_level": "anonymous",
    "port": 389,
    "preferred_ad_servers": {
    },
    "servers": {
    },
    "session_security": "none",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "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
23724130	Cannot use an IPv6 name server address because there are no IPv6 LIFs

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	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ldap\_service

Name	Type	Description
_links	<a href="#">_links</a>	
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.

Name	Type	Description
base_scope	string	Specifies the default search scope for LDAP queries: <ul style="list-style-type: none"> <li>• base - search the named entry only</li> <li>• onelevel - search all entries immediately below the DN</li> <li>• subtree - search the named DN entry and the entire subtree below the DN</li> </ul>
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"> <li>• anonymous - anonymous bind</li> <li>• simple - simple bind</li> <li>• sasl - Simple Authentication and Security Layer (SASL) bind</li> </ul>
port	integer	The port used to connect to the LDAP Servers.
preferred_ad_servers	array[string]	
schema	string	The name of the schema template used by the SVM. <ul style="list-style-type: none"> <li>• AD-IDMU - Active Directory Identity Management for UNIX</li> <li>• AD-SFU - Active Directory Services for UNIX</li> <li>• MS-AD-BIS - Active Directory Identity Management for UNIX</li> <li>• RFC-2307 - Schema based on RFC 2307</li> <li>• Custom schema</li> </ul>

Name	Type	Description
servers	array[string]	
session_security	string	Specifies the level of security to be used for LDAP communications: <ul style="list-style-type: none"> <li>• none - no signing or sealing</li> <li>• sign - sign LDAP traffic</li> <li>• seal - seal and sign LDAP traffic</li> </ul>
svm	<a href="#">svm</a>	
use_start_tls	boolean	Specifies whether or not to use Start TLS over LDAP connections.

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete the LDAP configuration for an SVM

```
DELETE /name-services/ldap/{svm.uuid}
```

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.

### Learn more

- [DOC /name-services/ldap](#)

### Parameters

Name	Type	In	Required	Description
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	<a href="#">error</a>	

### 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve the LDAP configuration for an SVM

GET /name-services/ldap/{svm.uuid}

Retrieves LDAP configuration for an SVM. All parameters for the LDAP configuration are displayed by default.

### Learn more

- [DOC /name-services/ldap](#)

### 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
<a href="#">_links</a>	<a href="#">_links</a>	
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"><li>• base - search the named entry only</li><li>• onelevel - search all entries immediately below the DN</li><li>• subtree - search the named DN entry and the entire subtree below the DN</li></ul>
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"><li>• anonymous - anonymous bind</li><li>• simple - simple bind</li><li>• sasl - Simple Authentication and Security Layer (SASL) bind</li></ul>
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"> <li>• AD-IDMU - Active Directory Identity Management for UNIX</li> <li>• AD-SFU - Active Directory Services for UNIX</li> <li>• MS-AD-BIS - Active Directory Identity Management for UNIX</li> <li>• RFC-2307 - Schema based on RFC 2307</li> <li>• Custom schema</li> </ul>
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"> <li>• none - no signing or sealing</li> <li>• sign - sign LDAP traffic</li> <li>• seal - seal and sign LDAP traffic</li> </ul>
svm	svm	
use_start_tls	boolean	<p>Specifies whether or not to use Start TLS over LDAP connections.</p>

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "base_scope": "base",
  "min_bind_level": "anonymous",
  "port": 389,
  "preferred_ad_servers": {
  },
  "servers": {
  },
  "session_security": "none",
  "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	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the LDAP configuration for an SVM

```
PATCH /name-services/ldap/{svm.uuid}
```

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:

1. The server does not have LDAP installed.
2. The server or Active Directory domain is invalid.
3. The server or Active Directory domain is unreachable

#### Learn more

- [DOC /name-services/ldap](#)

#### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

#### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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.

Name	Type	Description
base_scope	string	Specifies the default search scope for LDAP queries: <ul style="list-style-type: none"> <li>• base - search the named entry only</li> <li>• onelevel - search all entries immediately below the DN</li> <li>• subtree - search the named DN entry and the entire subtree below the DN</li> </ul>
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"> <li>• anonymous - anonymous bind</li> <li>• simple - simple bind</li> <li>• sasl - Simple Authentication and Security Layer (SASL) bind</li> </ul>
port	integer	The port used to connect to the LDAP Servers.
preferred_ad_servers	array[string]	
schema	string	The name of the schema template used by the SVM. <ul style="list-style-type: none"> <li>• AD-IDMU - Active Directory Identity Management for UNIX</li> <li>• AD-SFU - Active Directory Services for UNIX</li> <li>• MS-AD-BIS - Active Directory Identity Management for UNIX</li> <li>• RFC-2307 - Schema based on RFC 2307</li> <li>• Custom schema</li> </ul>
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"> <li>• none - no signing or sealing</li> <li>• sign - sign LDAP traffic</li> <li>• seal - seal and sign LDAP traffic</li> </ul>
svm	svm	
use_start_tls	boolean	Specifies whether or not to use Start TLS over LDAP connections.

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "base_scope": "base",
  "min_bind_level": "anonymous",
  "port": 389,
  "preferred_ad_servers": {
  },
  "servers": {
  },
  "session_security": "none",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "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
23724130	Cannot use an IPv6 name server address because there are no IPv6 LIFs

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	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ldap\_service

Name	Type	Description
_links	<a href="#">_links</a>	
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.

Name	Type	Description
base_scope	string	Specifies the default search scope for LDAP queries: <ul style="list-style-type: none"> <li>• base - search the named entry only</li> <li>• onelevel - search all entries immediately below the DN</li> <li>• subtree - search the named DN entry and the entire subtree below the DN</li> </ul>
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"> <li>• anonymous - anonymous bind</li> <li>• simple - simple bind</li> <li>• sasl - Simple Authentication and Security Layer (SASL) bind</li> </ul>
port	integer	The port used to connect to the LDAP Servers.
preferred_ad_servers	array[string]	
schema	string	The name of the schema template used by the SVM. <ul style="list-style-type: none"> <li>• AD-IDMU - Active Directory Identity Management for UNIX</li> <li>• AD-SFU - Active Directory Services for UNIX</li> <li>• MS-AD-BIS - Active Directory Identity Management for UNIX</li> <li>• RFC-2307 - Schema based on RFC 2307</li> <li>• Custom schema</li> </ul>

Name	Type	Description
servers	array[string]	
session_security	string	Specifies the level of security to be used for LDAP communications: <ul style="list-style-type: none"> <li>• none - no signing or sealing</li> <li>• sign - sign LDAP traffic</li> <li>• seal - seal and sign LDAP traffic</li> </ul>
svm	<a href="#">svm</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage name mappings for SVMs

### 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\":
\"ENGCIFS_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": "ENGCIFS_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\":
\"ENGCIFS_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": "ENGCIFS_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": "ENGCIFS_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\": \"ENGCIFS_LOCAL_USER\", \"replacement\": \"pcuser\"}"
```

---

## 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"
```

## Retrieve hostname mapping for all SVMs

GET /name-services/name-mappings

Retrieves the name mapping configuration for all SVMs.

### Related ONTAP commands

- `vserver name-mapping show`

### Learn more

- [DOC /name-services/name-mappings](#)

### Parameters

Name	Type	In	Required	Description
direction	string	query	False	Filter by direction
client_match	string	query	False	Filter by client_match
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
replacement	string	query	False	Filter by replacement
index	integer	query	False	Filter by index
pattern	string	query	False	Filter by pattern
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">name_mapping</a> ]	

## Example response

```
{
  "_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"
    }
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
client_match	string	<p>Client workstation IP Address which is matched when searching for the pattern.</p> <p>You can specify the value in any of the following formats:</p> <ul style="list-style-type: none"> <li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li> <li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li> <li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li> <li>• As a hostname</li> </ul>
direction	string	<p>Direction in which the name mapping is applied. The possible values are:</p> <ul style="list-style-type: none"> <li>• krb_unix - Kerberos principal name to UNIX user name</li> <li>• win_unix - Windows user name to UNIX user name</li> <li>• unix_win - UNIX user name to Windows user name mapping</li> </ul>
index	integer	<p>Position in the list of name mappings.</p>
pattern	string	<p>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.</p>



Name	Type	Description
replacement	string	The name that is used as a replacement, if the pattern associated with this entry matches.
svm	<a href="#">svm</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create the hostname mappings for an SVM

POST /name-services/name-mappings

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

- [DOC /name-services/name-mappings](#)

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>client_match</code>	string	<p>Client workstation IP Address which is matched when searching for the pattern.</p> <p>You can specify the value in any of the following formats:</p> <ul style="list-style-type: none"><li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li><li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li><li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li><li>• As a hostname</li></ul>
<code>direction</code>	string	<p>Direction in which the name mapping is applied. The possible values are:</p> <ul style="list-style-type: none"><li>• <code>krb_unix</code> - Kerberos principal name to UNIX user name</li><li>• <code>win_unix</code> - Windows user name to UNIX user name</li><li>• <code>unix_win</code> - UNIX user name to Windows user name mapping</li></ul>
<code>index</code>	integer	Position in the list of name mappings.

Name	Type	Description
pattern	string	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	The name that is used as a replacement, if the pattern associated with this entry matches.
svm	svm	

### Example request

```
{
  "_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"
  }
}
```

### Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">name_mapping</a> ]	

### Example response

```
{
  "_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": "ENGCIFS_AD_USER",
    "replacement": "unix_user1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "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

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	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
client_match	string	<p>Client workstation IP Address which is matched when searching for the pattern.</p> <p>You can specify the value in any of the following formats:</p> <ul style="list-style-type: none"> <li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li> <li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li> <li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li> <li>• As a hostname</li> </ul>
direction	string	<p>Direction in which the name mapping is applied. The possible values are:</p> <ul style="list-style-type: none"> <li>• krb_unix - Kerberos principal name to UNIX user name</li> <li>• win_unix - Windows user name to UNIX user name</li> <li>• unix_win - UNIX user name to Windows user name mapping</li> </ul>
index	integer	<p>Position in the list of name mappings.</p>
pattern	string	<p>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.</p>

Name	Type	Description
replacement	string	The name that is used as a replacement, if the pattern associated with this entry matches.
svm	<a href="#">svm</a>	

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete the name mapping configuration

DELETE /name-services/name-mappings/{svm.uuid}/{direction}/{index}

Deletes the name mapping configuration.

#### Related ONTAP commands

- `vserver name-mapping delete`



## Learn more

- [DOC /name-services/name-mappings](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
direction	string	path	True	Direction
index	integer	path	True	Position of the entry in the list

## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve the name mapping configuration for an SVM

GET /name-services/name-mappings/{svm.uuid}/{direction}/{index}

Retrieves the name mapping configuration of an SVM.

### Related ONTAP commands

- `vserver name-mapping show`

### Learn more

- [DOC /name-services/name-mappings](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
direction	string	path	True	Direction

Name	Type	In	Required	Description
index	integer	path	True	Position of the entry in the list
fields	array[string]	query	False	Specify the fields to return.

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
client_match	string	<p>Client workstation IP Address which is matched when searching for the pattern. You can specify the value in any of the following formats:</p> <ul style="list-style-type: none"> <li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li> <li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li> <li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li> <li>• As a hostname</li> </ul>
direction	string	<p>Direction in which the name mapping is applied. The possible values are:</p> <ul style="list-style-type: none"> <li>• krb_unix - Kerberos principal name to UNIX user name</li> <li>• win_unix - Windows user name to UNIX user name</li> <li>• unix_win - UNIX user name to Windows user name mapping</li> </ul>
index	integer	Position in the list of name mappings.

Name	Type	Description
pattern	string	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	The name that is used as a replacement, if the pattern associated with this entry matches.
svm	svm	

### Example response

```
{
  "_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

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	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the name mapping configuration for an SVM

PATCH /name-services/name-mappings/{svm.uuid}/{direction}/{index}

Updates the name mapping configuration of an SVM.

#### Related ONTAP commands

- `vserver name-mapping insert`
- `vserver name-mapping modify`

#### Learn more

- [DOC /name-services/name-mappings](#)

#### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
direction	string	path	True	Direction
index	integer	path	True	Position of the entry in the list

#### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
client_match	string	<p>Client workstation IP Address which is matched when searching for the pattern.</p> <p>You can specify the value in any of the following formats:</p> <ul style="list-style-type: none"><li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li><li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li><li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li><li>• As a hostname</li></ul>

Name	Type	Description
direction	string	Direction in which the name mapping is applied. The possible values are: <ul style="list-style-type: none"> <li>• krb_unix - Kerberos principal name to UNIX user name</li> <li>• win_unix - Windows user name to UNIX user name</li> <li>• unix_win - UNIX user name to Windows user name mapping</li> </ul>
index	integer	Position in the list of name mappings.
pattern	string	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	The name that is used as a replacement, if the pattern associated with this entry matches.
svm	svm	



## Example request

```
{
  "_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"
  }
}
```

## Response

Status: 200, Ok

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
65798185	Failed to resolve the specified hostname

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	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
client_match	string	<p>Client workstation IP Address which is matched when searching for the pattern.</p> <p>You can specify the value in any of the following formats:</p> <ul style="list-style-type: none"> <li>• As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24</li> <li>• As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64</li> <li>• As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0</li> <li>• As a hostname</li> </ul>
direction	string	<p>Direction in which the name mapping is applied. The possible values are:</p> <ul style="list-style-type: none"> <li>• krb_unix - Kerberos principal name to UNIX user name</li> <li>• win_unix - Windows user name to UNIX user name</li> <li>• unix_win - UNIX user name to Windows user name mapping</li> </ul>
index	integer	<p>Position in the list of name mappings.</p>
pattern	string	<p>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.</p>

Name	Type	Description
replacement	string	The name that is used as a replacement, if the pattern associated with this entry matches.
svm	<a href="#">svm</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage NIS server configurations

### 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 </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"
```

## Retrieve NIS domain configurations of all SVMs

GET /name-services/nis

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

- [DOC /name-services/nis](#)

### 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	The default is true for GET calls. When set to false, only the number of records is returned.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of NIS domain records.
records	array[ <a href="#">nis_service</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "bound_servers": {
    },
    "servers": {
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  }
}
```

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

nis\_service

Name	Type	Description
_links	<a href="#">_links</a>	
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.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create the NIS domain and server configuration for a data SVM

POST `/name-services/nis`

Creates an NIS domain and server configuration for a data SVM.

NIS configuration for the cluster is managed via [/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

- [DOC /name-services/nis](#)

## Request Body

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>bound_servers</code>	<code>array[string]</code>	
<code>domain</code>	<code>string</code>	The NIS domain to which this configuration belongs.
<code>servers</code>	<code>array[string]</code>	A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration.
<code>svm</code>	<code>svm</code>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bound_servers": {
  },
  "servers": {
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
```

## Response

Status: 201, Created

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>num_records</code>	integer	Number of NIS domain records.
<code>records</code>	array[nis_service]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "bound_servers": {
    },
    "servers": {
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "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
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

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

nis\_service

Name	Type	Description
_links	<a href="#">_links</a>	
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.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete the NIS domain configuration for an SVM

```
DELETE /name-services/nis/{svm.uuid}
```

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

- [DOC /name-services/nis](#)

#### Parameters

Name	Type	In	Required	Description
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve the NIS domain and server configurations for an SVM

```
GET /name-services/nis/{svm.uuid}
```

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

- [DOC /name-services/nis](#)

### Parameters

Name	Type	In	Required	Description
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
<a href="#">_links</a>	<a href="#">_links</a>	
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.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bound_servers": {
  },
  "servers": {
  },
  "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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the NIS domain and server configuration for an SVM

PATCH /name-services/nis/{svm.uuid}

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

- [DOC /name-services/nis](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bound_servers": {
  },
  "servers": {
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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

Error Code	Description
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

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

nis\_service

Name	Type	Description
_links	<a href="#">_links</a>	
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.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument



error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Networking

### 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

### Network Ethernet broadcast-domains endpoint overview

#### Overview

A broadcast domain is a collection of Ethernet ports that have layer2 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",
  "ipospace": {
    "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": "ace1a36f-4bed-11e9-af6a-005056bb13c0",
          "name": "e0b",
```

```
    "node": {
      "name": "examplecluster-node-1"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/ace1a36f-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

The broadcast domains POST API is used 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

The broadcast domain PATCH API is used to update attributes of broadcast domains.

---

#### Example

#### 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 }'
{
}
```

---

### Deleting network Ethernet broadcast domains

The broadcast domain DELETE API is used 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/"
```

---

### Retrieve broadcast domains for the entire cluster

GET /network/ethernet/broadcast-domains

Retrieves a collection of broadcast domains for the entire cluster.

#### Related ONTAP commands

- `network port broadcast-domain show`

#### Learn more

- [DOC /network/ethernet/broadcast-domains](#)

## Parameters

Name	Type	In	Required	Description
ports.node.name	string	query	False	Filter by ports.node.name
ports.name	string	query	False	Filter by ports.name
ports.uuid	string	query	False	Filter by ports.uuid
name	string	query	False	Filter by name
mtu	integer	query	False	Filter by mtu
uuid	string	query	False	Filter by uuid
ipspace.uuid	string	query	False	Filter by ipspace.uuid
ipspace.name	string	query	False	Filter by ipspace.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.
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.

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	<a href="#">_links</a>	
num_records	integer	
records	array[ <a href="#">broadcast_domain</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_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

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

ipSPACE

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

node

Name	Type	Description
name	string	Name of node on which the port is located.

ports

Port UUID along with readable names

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

## broadcast\_domain

Set of ports that will receive a broadcast Ethernet packet from any of them

Name	Type	Description
_links	<a href="#">_links</a>	
ipspace	<a href="#">ipspace</a>	Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.
mtu	integer	Maximum transmission unit, largest packet size on this network
name	string	Name of the broadcast domain, scoped to its IPspace
ports	array[ <a href="#">ports</a> ]	Ports that belong to the broadcast domain
uuid	string	Broadcast domain UUID

## error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

## error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create a new broadcast domain

POST /network/ethernet/broadcast-domains

Creates a new broadcast domain.

#### Required properties

- `name` - Name of the broadcast-domain to create.
- `mtu` - Maximum transmission unit 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`

#### Learn more

- [DOC /network/ethernet/broadcast-domains](#)

#### Request Body

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>ipSpace</code>	<code>ipSpace</code>	Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.
<code>mtu</code>	integer	Maximum transmission unit, largest packet size on this network
<code>name</code>	string	Name of the broadcast domain, scoped to its IPspace
<code>ports</code>	array[ <code>ports</code> ]	Ports that belong to the broadcast domain
<code>uuid</code>	string	Broadcast domain UUID

## Example request

```
{
  "_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": "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
1967082	The specified ipspace.name does not match the IPspace name of ipspace.uuid.
1967102	POST operation might have left the configuration in an inconsistent state. Check the configuration.
1377267	The specified IPspace 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

ipSPACE

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

node

Name	Type	Description
name	string	Name of node on which the port is located.

ports

Port UUID along with readable names

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

broadcast\_domain

Set of ports that will receive a broadcast Ethernet packet from any of them

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
ipspace	<a href="#">ipspace</a>	Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.
mtu	integer	Maximum transmission unit, largest packet size on this network
name	string	Name of the broadcast domain, scoped to its IPspace
ports	array[ <a href="#">ports</a> ]	Ports that belong to the broadcast domain
uuid	string	Broadcast domain UUID

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete a broadcast domain

```
DELETE /network/ethernet/broadcast-domains/{uuid}
```

Deletes a broadcast domain.

## Related ONTAP commands

- `network port broadcast-domain delete`

## Learn more

- [DOC /network/ethernet/broadcast-domains](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

## Response

```
Status: 200, Ok
```

## Error

```
Status: Default
```

## ONTAP Error Response Codes

Error Code	Description
1967103	A broadcast domain with ports cannot be deleted.

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve broadcast domain details

GET /network/ethernet/broadcast-domains/{uuid}

Retrieves details of a broadcast domain.

### Related ONTAP commands

- `network port broadcast-domain show`

### Learn more

- [DOC /network/ethernet/broadcast-domains](#)

### Parameters

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

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
ipspace	<a href="#">ipspace</a>	Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.
mtu	integer	Maximum transmission unit, largest packet size on this network
name	string	Name of the broadcast domain, scoped to its IPspace
ports	array[ <a href="#">ports</a> ]	Ports that belong to the broadcast domain
uuid	string	Broadcast domain UUID

## Example response

```
{
  "_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": "e1b",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "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	<a href="#">href</a>	

ipSPACE

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

node

Name	Type	Description
name	string	Name of node on which the port is located.

ports

Port UUID along with readable names

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

error\_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update broadcast domain properties

PATCH /network/ethernet/broadcast-domains/{uuid}

Updates the properties of a broadcast domain.

### Related ONTAP commands

- `network port broadcast-domain modify`
- `network port broadcast-domain rename`

### Learn more

- [DOC /network/ethernet/broadcast-domains](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Broadcast domain UUID

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
ipspace	<a href="#">ipspace</a>	Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

Name	Type	Description
mtu	integer	Maximum transmission unit, largest packet size on this network
name	string	Name of the broadcast domain, scoped to its IPspace
ports	array[ports]	Ports that belong to the broadcast domain
uuid	string	Broadcast domain UUID

## Example request

```
{
  "_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": "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
1967082	The specified ipspace.name does not match the IPspace name of ipspace.uuid.
1377267	The specified IPspace 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

node

Name	Type	Description
name	string	Name of node on which the port is located.

ports

Port UUID along with readable names

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

broadcast\_domain

Set of ports that will receive a broadcast Ethernet packet from any of them

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
ipspace	<a href="#">ipspace</a>	Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.
mtu	integer	Maximum transmission unit, largest packet size on this network
name	string	Name of the broadcast domain, scoped to its IPspace
ports	array[ <a href="#">ports</a> ]	Ports that belong to the broadcast domain
uuid	string	Broadcast domain UUID

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage network Ethernet ports

### 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).

The GET (collection), GET (instance) and PATCH APIs are available for all port types. The 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",
  "_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"
      }
    }
  }
}

```



```

    }
  },
  "enabled": true,
  "state": "up",
  "mtu": 1500,
  "speed": "1000",
  "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"
      }
    }
  },
  "enabled": true,
  "state": "up",
  "mtu": 1500,
  "speed": "1000",
  "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

The network ports POST API is used to create VLAN and LAG ports.

## 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". 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\" },  
\"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:netapp! -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

The network ports PATCH API is used to modify 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\" }"
```

---

## Deleting ports

The network ports DELETE API is used 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"
```

## Retrieve ports

GET /network/ethernet/ports

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`

### Learn more

- [DOC /network/ethernet/ports](#)

### Parameters

Name	Type	In	Required	Description
name	string	query	False	Filter by name
state	string	query	False	Filter by state
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
broadcast_domain.name	string	query	False	Filter by broadcast_domain.name
broadcast_domain.ip space.name	string	query	False	Filter by broadcast_domain.ip space.name

Name	Type	In	Required	Description
broadcast_domain.uuid	string	query	False	Filter by broadcast_domain.uuid
speed	integer	query	False	Filter by speed
enabled	boolean	query	False	Filter by enabled
lag.active_ports.node.name	string	query	False	Filter by lag.active_ports.node.name
lag.active_ports.name	string	query	False	Filter by lag.active_ports.name
lag.active_ports.uuid	string	query	False	Filter by lag.active_ports.uuid
lag.distribution_policy	string	query	False	Filter by lag.distribution_policy
lag.mode	string	query	False	Filter by lag.mode
lag.member_ports.node.name	string	query	False	Filter by lag.member_ports.node.name
lag.member_ports.name	string	query	False	Filter by lag.member_ports.name
lag.member_ports.uuid	string	query	False	Filter by lag.member_ports.uuid
mtu	integer	query	False	Filter by mtu
uuid	string	query	False	Filter by uuid
mac_address	string	query	False	Filter by mac_address

Name	Type	In	Required	Description
vlan.base_port.node.name	string	query	False	Filter by vlan.base_port.node.name
vlan.base_port.name	string	query	False	Filter by vlan.base_port.name
vlan.base_port.uuid	string	query	False	Filter by vlan.base_port.uuid
vlan.tag	integer	query	False	Filter by vlan.tag
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	
records	array[ <a href="#">port</a> ]	

## Example response

```
{
  "_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": "port",
      "member_ports": {
        "_links": {
          "self": {
```

```

        "href": "/api/resourcelink"
    }
},
"name": "elb",
"node": {
    "name": "node1"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"mode": "multimode_lacp"
},
"mac_address": "01:02:03:04:05:06",
"mtu": 1500,
"name": "elb",
"node": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"speed": 1000,
"state": "up",
"type": "vlan",
"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

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

ipospace

Name	Type	Description
name	string	Name of the broadcast domain's IPspace

broadcast\_domain

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

Name	Type	Description
_links	<a href="#">_links</a>	
ipospace	<a href="#">ipospace</a>	
name	string	Name of the broadcast domain, scoped to its IPspace
uuid	string	Broadcast domain UUID

node

Name	Type	Description
name	string	Name of node on which the port is located.

## active\_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

## member\_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

## lag

Name	Type	Description
active_ports	array[ <a href="#">active_ports</a> ]	Active ports of a LAG (ifgrp). (Some member ports may be inactive.)
distribution_policy	string	Policy for mapping flows to ports for outbound packets through a LAG (ifgrp).
member_ports	array[ <a href="#">member_ports</a> ]	
mode	string	Determines how the ports interact with the switch.

## node

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

## base\_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

vlan

Name	Type	Description
base_port	<a href="#">base_port</a>	Port UUID along with readable names. Either the UUID or both names may be supplied on input.
tag	integer	VLAN ID

port

Name	Type	Description
_links	<a href="#">_links</a>	
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.
enabled	boolean	
lag	<a href="#">lag</a>	
mac_address	string	
mtu	integer	MTU of the port in bytes. Set by broadcast domain.
name	string	Portname, such as e0a, e1b-100 (VLAN on ethernet), a0c (LAG/ifgrp), a0d-200 (vlan on LAG/ifgrp)
node	<a href="#">node</a>	
speed	integer	Link speed in Mbps
state	string	Operational state of the port.
type	string	Type of physical or virtual port

Name	Type	Description
uuid	string	Port UUID
vlan	<a href="#">vlan</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create a new VLAN or LAG

POST /network/ethernet/ports

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.
- `broadcast_domain` - Broadcast domain the port is associated with.
- `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.

### Related ONTAP commands

- `network port ifgrp create`
- `network port vlan create`

### Learn more

- [DOC /network/ethernet/ports](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>broadcast_domain</code>	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.
<code>enabled</code>	boolean	
<code>lag</code>	<a href="#">lag</a>	
<code>mac_address</code>	string	
<code>mtu</code>	integer	MTU of the port in bytes. Set by broadcast domain.
<code>name</code>	string	Portname, such as e0a, e1b-100 (VLAN on ethernet), a0c (LAG/ifgrp), a0d-200 (vlan on LAG/ifgrp)
<code>node</code>	<a href="#">node</a>	
<code>speed</code>	integer	Link speed in Mbps
<code>state</code>	string	Operational state of the port.
<code>type</code>	string	Type of physical or virtual port
<code>uuid</code>	string	Port UUID
<code>vlan</code>	<a href="#">vlan</a>	

## Example request

```
{
  "_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": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "distribution_policy": "port",
    "member_ports": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
}
```

```

    "mode": "multimode_lacp"
  },
  "mac_address": "01:02:03:04:05:06",
  "mtu": 1500,
  "name": "e1b",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "speed": 1000,
  "state": "up",
  "type": "vlan",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vlan": {
    "base_port": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    },
    "name": "e1b",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "tag": 100
}
}

```

## Response

Status: 201, Created

Name	Type	Description
_links	_links	
num_records	integer	
records	array[port]	

## Example response

```
{
  "_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": "port",
      "member_ports": {
        "_links": {
          "self": {
```



```

        "href": "/api/resourcelink"
    }
},
"name": "elb",
"node": {
    "name": "node1"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"mode": "multimode_lacp"
},
"mac_address": "01:02:03:04:05:06",
"mtu": 1500,
"name": "elb",
"node": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"speed": 1000,
"state": "up",
"type": "vlan",
"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

### ONTAP Error Response Codes

Error Code	Description
1967083	The specified type is invalid.
1967084	The specified node UUID is invalid.
1967085	The specified node name is invalid.
1967086	Node name and UUID must match if both are provided.
1967087	The specified broadcast domain UUID is invalid.
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 invalid.
1967091	The specified VLAN base port name and node name is invalid.
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 invalid.
1967095	The specified LAG member port name and node name combination is invalid.
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.

Error Code	Description
1967148	Failure to remove port from broadcast domain.
1967149	Failure to add port to broadcast domain.
1376361	Port is already a member of a LAG (ifgrp).
1966189	Port is the home port or current port of a LIF.

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	<a href="#">href</a>	

ipospace

Name	Type	Description
name	string	Name of the broadcast domain's IPspace

broadcast\_domain

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

Name	Type	Description
_links	<a href="#">_links</a>	
ipospace	<a href="#">ipospace</a>	
name	string	Name of the broadcast domain, scoped to its IPspace
uuid	string	Broadcast domain UUID

node

Name	Type	Description
name	string	Name of node on which the port is located.

active\_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	
node	<a href="#">node</a>	
uuid	string	

member\_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

lag

Name	Type	Description
active_ports	array[ <a href="#">active_ports</a> ]	Active ports of a LAG (ifgrp). (Some member ports may be inactive.)
distribution_policy	string	Policy for mapping flows to ports for outbound packets through a LAG (ifgrp).
member_ports	array[ <a href="#">member_ports</a> ]	
mode	string	Determines how the ports interact with the switch.

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

base\_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	
node	<a href="#">node</a>	
uuid	string	

vlan

Name	Type	Description
base_port	<a href="#">base_port</a>	Port UUID along with readable names. Either the UUID or both names may be supplied on input.
tag	integer	VLAN ID

port

Name	Type	Description
_links	<a href="#">_links</a>	
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.
enabled	boolean	
lag	<a href="#">lag</a>	
mac_address	string	
mtu	integer	MTU of the port in bytes. Set by broadcast domain.
name	string	Portname, such as e0a, e1b-100 (VLAN on ethernet), a0c (LAG/ifgrp), a0d-200 (vlan on LAG/ifgrp)
node	<a href="#">node</a>	
speed	integer	Link speed in Mbps
state	string	Operational state of the port.
type	string	Type of physical or virtual port
uuid	string	Port UUID

Name	Type	Description
vlan	<a href="#">vlan</a>	

[\\_links](#)

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

[error\\_arguments](#)

Name	Type	Description
code	string	Argument code
message	string	Message argument

[error](#)

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete a VLAN or LAG

DELETE /network/ethernet/ports/{uuid}

Deletes a VLAN or LAG (ifgrp).

### Related ONTAP commands

- `network port ifgrp delete`
- `network port vlan delete`

### Learn more

- [DOC /network/ethernet/ports](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

## Response

Status: 200, Ok

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
1966189	Port is the home port or current port of a LIF.
1376858	Port already has a LIF bound.

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve a physical port, VLAN, or LAG details

GET /network/ethernet/ports/{uuid}

Retrieves the details of a physical port, VLAN, or LAG.

### Related ONTAP commands

- `network port show`
- `network port ifgrp show`
- `network port vlan show`

### Learn more

- [DOC /network/ethernet/ports](#)

### Parameters

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

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.
enabled	boolean	
lag	<a href="#">lag</a>	
mac_address	string	
mtu	integer	MTU of the port in bytes. Set by broadcast domain.
name	string	Portname, such as e0a, e1b-100 (VLAN on ethernet), a0c (LAG/ifgrp), a0d-200 (vlan on LAG/ifgrp)
node	<a href="#">node</a>	
speed	integer	Link speed in Mbps
state	string	Operational state of the port.
type	string	Type of physical or virtual port
uuid	string	Port UUID
vlan	<a href="#">vlan</a>	

## Example response

```
{
  "_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": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "distribution_policy": "port",
    "member_ports": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
}
```

```

    "mode": "multimode_lacp"
  },
  "mac_address": "01:02:03:04:05:06",
  "mtu": 1500,
  "name": "e1b",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "speed": 1000,
  "state": "up",
  "type": "vlan",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vlan": {
    "base_port": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    },
    "name": "e1b",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "tag": 100
}
}

```

## 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	<a href="#">href</a>	

ipospace

Name	Type	Description
name	string	Name of the broadcast domain's IPspace

broadcast\_domain

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

Name	Type	Description
_links	<a href="#">_links</a>	
ipospace	<a href="#">ipospace</a>	
name	string	Name of the broadcast domain, scoped to its IPspace
uuid	string	Broadcast domain UUID

node

Name	Type	Description
name	string	Name of node on which the port is located.

active\_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	
node	<a href="#">node</a>	
uuid	string	

member\_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

lag

Name	Type	Description
active_ports	array[ <a href="#">active_ports</a> ]	Active ports of a LAG (ifgrp). (Some member ports may be inactive.)
distribution_policy	string	Policy for mapping flows to ports for outbound packets through a LAG (ifgrp).
member_ports	array[ <a href="#">member_ports</a> ]	
mode	string	Determines how the ports interact with the switch.

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

base\_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	
node	<a href="#">node</a>	
uuid	string	

vlan

Name	Type	Description
base_port	<a href="#">base_port</a>	Port UUID along with readable names. Either the UUID or both names may be supplied on input.
tag	integer	VLAN ID

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update a port

PATCH /network/ethernet/ports/{uuid}

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

#### Learn more

- [DOC /network/ethernet/ports](#)

#### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Port UUID

#### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.
enabled	boolean	
lag	<a href="#">lag</a>	
mac_address	string	
mtu	integer	MTU of the port in bytes. Set by broadcast domain.
name	string	Portname, such as e0a, e1b-100 (VLAN on ethernet), a0c (LAG/ifgrp), a0d-200 (vlan on LAG/ifgrp)
node	<a href="#">node</a>	
speed	integer	Link speed in Mbps
state	string	Operational state of the port.
type	string	Type of physical or virtual port
uuid	string	Port UUID
vlan	<a href="#">vlan</a>	

## Example request

```
{
  "_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": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "distribution_policy": "port",
    "member_ports": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
}
```

```
    "mode": "multimode_lacp",
  },
  "mac_address": "01:02:03:04:05:06",
  "mtu": 1500,
  "name": "e1b",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "speed": 1000,
  "state": "up",
  "type": "vlan",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vlan": {
    "base_port": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    },
    "name": "e1b",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "tag": 100
}
}
```

#### Response

Status: 200, Ok

#### Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
1967087	The specified broadcast domain UUID is invalid.
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 invalid.
1967095	The specified LAG member port name and node name combination is invalid.
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.
1376361	Port is already a member of a LAG (ifgrp).
1377563	Port is already a member of a LAG (ifgrp).
1377562	Port cannot be used because it is currently the home port or current port of a LIF.

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	<a href="#">href</a>	

ipospace

Name	Type	Description
name	string	Name of the broadcast domain's IPspace

broadcast\_domain

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

Name	Type	Description
_links	<a href="#">_links</a>	
ipospace	<a href="#">ipospace</a>	
name	string	Name of the broadcast domain, scoped to its IPspace
uuid	string	Broadcast domain UUID

node

Name	Type	Description
name	string	Name of node on which the port is located.

active\_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	
node	<a href="#">node</a>	
uuid	string	

member\_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

lag

Name	Type	Description
active_ports	array[ <a href="#">active_ports</a> ]	Active ports of a LAG (ifgrp). (Some member ports may be inactive.)
distribution_policy	string	Policy for mapping flows to ports for outbound packets through a LAG (ifgrp).
member_ports	array[ <a href="#">member_ports</a> ]	
mode	string	Determines how the ports interact with the switch.

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

base\_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	
node	<a href="#">node</a>	
uuid	string	

vlan

Name	Type	Description
base_port	<a href="#">base_port</a>	Port UUID along with readable names. Either the UUID or both names may be supplied on input.
tag	integer	VLAN ID

port

Name	Type	Description
_links	<a href="#">_links</a>	
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.
enabled	boolean	
lag	<a href="#">lag</a>	
mac_address	string	
mtu	integer	MTU of the port in bytes. Set by broadcast domain.
name	string	Portname, such as e0a, e1b-100 (VLAN on ethernet), a0c (LAG/ifgrp), a0d-200 (vlan on LAG/ifgrp)
node	<a href="#">node</a>	
speed	integer	Link speed in Mbps
state	string	Operational state of the port.
type	string	Type of physical or virtual port
uuid	string	Port UUID



Name	Type	Description
vlan	vlan	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	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 FC network interfaces

### 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.

#### 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": { "port": { "name": "0a", "node": { "name": "node1" } } },
"data_protocol": "fcp" }'

# 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": {
        "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": "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"
  }
}
]
}

```

### 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": { "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": {
    "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": {
        "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": {
        "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"
      }
    }
  ],
  "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": {
    "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"
  }
}
}
```

### 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 the 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": { "port": { "uuid": "a1dc7aa5-db83-11e8-9ef7-005056bbbbcc" } }
}'
```

## Deleting an FC interface

An 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'
```

## Retrieve FC interfaces

GET /network/fc/interfaces

Retrieves FC interfaces.

### Related ONTAP commands

- `network interface show`
- `vserver fcp interface show`

### Learn more

- [DOC /network/fc/interfaces](#)

### Parameters

Name	Type	In	Required	Description
wwpn	string	query	False	Filter by wwpn
uuid	string	query	False	Filter by uuid
comment	string	query	False	Filter by comment

Name	Type	In	Required	Description
location.port.name	string	query	False	Filter by location.port.name
location.port.node.name	string	query	False	Filter by location.port.node.name
location.port.uuid	string	query	False	Filter by location.port.uuid
location.node.name	string	query	False	Filter by location.node.name
location.node.uuid	string	query	False	Filter by location.node.uuid
port_address	string	query	False	Filter by port_address
wwnn	string	query	False	Filter by wwnn
enabled	boolean	query	False	Filter by enabled
state	string	query	False	Filter by state
name	string	query	False	Filter by name
data_protocol	string	query	False	Filter by data_protocol
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.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">fc_interface</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "comment": "string",
    "data_protocol": "fcp",
    "location": {
      "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"
      }
    },
    "name": "lif1",
    "port_address": "5060F",
    "state": "up",
    "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

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

node

The node on which the FC port is located.

Name	Type	Description
name	string	The name of the node on which the FC port is located.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FC port.
node	<a href="#">node</a>	The node on which the FC port is located.



Name	Type	Description
uuid	string	The unique identifier of the FC port.

#### 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.

Name	Type	Description
node	<a href="#">node</a>	
port	<a href="#">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.

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### 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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	A user configurable comment. Optional in POST; valid in PATCH. To clear a prior comment, set the property to an empty string in PATCH.

Name	Type	Description
data_protocol	string	The data protocol for which the FC interface is configured. Required in POST.
enabled	boolean	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 <i>true</i> (enabled) in POST.
location	<a href="#">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.
name	string	The name of the FC interface. Required in POST; optional in PATCH.
port_address	string	<p>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).</p> <p>This is useful for obtaining statistics and diagnostic information from FC switches.</p> <p>This is a hexadecimal encoded numeric value.</p>

Name	Type	Description
state	string	The current operational state of the FC interface. The state is set to <i>down</i> if the interface is not enabled.  If the node hosting the port is down or unavailable, no state value is returned.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the FC interface. Required in the URL.
wwnn	string	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.  <ul style="list-style-type: none"> <li>example: 20:00:00:50:56:b4:13:01</li> <li>readOnly: 1</li> </ul>
wwpn	string	The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created.  <ul style="list-style-type: none"> <li>example: 20:00:00:50:56:b4:13:a8</li> <li>readOnly: 1</li> </ul>

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create an FC interface

POST `/network/fc/interfaces`

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

- [DOC /network/fc/interfaces](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
comment	string	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	The data protocol for which the FC interface is configured. Required in POST.
enabled	boolean	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 <i>true</i> (enabled) in POST.
location	<a href="#">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.
name	string	The name of the FC interface. Required in POST; optional in PATCH.

Name	Type	Description
port_address	string	<p>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).</p> <p>This is useful for obtaining statistics and diagnostic information from FC switches.</p> <p>This is a hexadecimal encoded numeric value.</p>
state	string	<p>The current operational state of the FC interface. The state is set to <i>down</i> if the interface is not enabled.</p> <p>If the node hosting the port is down or unavailable, no state value is returned.</p>
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the FC interface. Required in the URL.
wwnn	string	<p>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.</p> <ul style="list-style-type: none"> <li>• example: 20:00:00:50:56:b4:13:01</li> <li>• readOnly: 1</li> </ul>

Name	Type	Description
wwpn	string	<p>The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created.</p> <ul style="list-style-type: none"><li>• example: 20:00:00:50:56:b4:13:a8</li><li>• readOnly: 1</li></ul>

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "data_protocol": "fcp",
  "location": {
    "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"
    }
  },
  "name": "lif1",
  "port_address": "5060F",
  "state": "up",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
}
```



```
"wvnn": "20:00:00:50:56:b4:13:01",  
"wvvn": "20:00:00:50:56:b4:13:a8"  
}
```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">fc_interface</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "comment": "string",
    "data_protocol": "fcp",
    "location": {
      "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"
      }
    },
    "name": "lif1",
    "port_address": "5060F",
    "state": "up",
    "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

Status: Default

## ONTAP Error Response Codes

Error Code	Description
1966140	A LIF 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 <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.
5374102	The specified FC LIF 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 <code>location.port.node.name</code> or <code>location.port.name</code> is supplied, both properties must be supplied.
5374873	The Fibre Channel port must be specified using either <code>location.port.uuid</code> or <code>location.port.node.name</code> and <code>location.port.name</code> .

Error Code	Description
72089652	An NVMe service must be created before creating a Fibre Channel interface using the NVMe over Fibre Channel data protocol.
72089672	The specified Fibre Channel port does not support the NVMe over Fibre Channel data protocol.
72089900	A Fibre Channel interface with the <i>fc_nvme</i> protocol cannot be creating in an SVM that is configured for a SAN protocol.
5373966	A Fibre Channel interface with the <i>fc</i> protocol cannot be creating in an SVM that is configured for NVMe.

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	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

node

The node on which the FC port is located.

Name	Type	Description
name	string	The name of the node on which the FC port is located.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FC port.
node	<a href="#">node</a>	The node on which the FC port is located.
uuid	string	The unique identifier of the FC port.

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.

Name	Type	Description
node	<a href="#">node</a>	
port	<a href="#">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.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	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	The data protocol for which the FC interface is configured. Required in POST.

Name	Type	Description
enabled	boolean	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 <i>true</i> (enabled) in POST.
location	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.
name	string	The name of the FC interface. Required in POST; optional in PATCH.
port_address	string	<p>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).</p> <p>This is useful for obtaining statistics and diagnostic information from FC switches.</p> <p>This is a hexadecimal encoded numeric value.</p>
state	string	<p>The current operational state of the FC interface. The state is set to <i>down</i> if the interface is not enabled.</p> <p>If the node hosting the port is down or unavailable, no state value is returned.</p>

Name	Type	Description
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the FC interface. Required in the URL.
wwnn	string	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. <ul style="list-style-type: none"> <li>example: 20:00:00:50:56:b4:13:01</li> <li>readOnly: 1</li> </ul>
wwpn	string	The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created. <ul style="list-style-type: none"> <li>example: 20:00:00:50:56:b4:13:a8</li> <li>readOnly: 1</li> </ul>

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments



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

### Delete an FC interface

DELETE /network/fc/interfaces/{uuid}

Deletes an FC interface.

#### Related ONTAP commands

- `network interface delete`

#### Learn more

- [DOC /network/fc/interfaces](#)

#### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

#### 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.

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an FC interface

GET /network/fc/interfaces/{uuid}

Retrieves a Fibre Channel interface.

## Related ONTAP commands

- `network interface show`
- `vserver fcp interface show`

## Learn more

- [DOC /network/fc/interfaces](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier for the FC interface.
fields	array[string]	query	False	Specify the fields to return.

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	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	The data protocol for which the FC interface is configured. Required in POST.
enabled	boolean	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 <i>true</i> (enabled) in POST.

Name	Type	Description
location	<a href="#">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.
name	string	The name of the FC interface. Required in POST; optional in PATCH.
port_address	string	<p>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).</p> <p>This is useful for obtaining statistics and diagnostic information from FC switches.</p> <p>This is a hexadecimal encoded numeric value.</p>
state	string	<p>The current operational state of the FC interface. The state is set to <i>down</i> if the interface is not enabled.</p> <p>If the node hosting the port is down or unavailable, no state value is returned.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the FC interface. Required in the URL.

Name	Type	Description
wwnn	string	<p>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.</p> <ul style="list-style-type: none"> <li>• example: 20:00:00:50:56:b4:13:01</li> <li>• readOnly: 1</li> </ul>
wwpn	string	<p>The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created.</p> <ul style="list-style-type: none"> <li>• example: 20:00:00:50:56:b4:13:a8</li> <li>• readOnly: 1</li> </ul>

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "data_protocol": "fcp",
  "location": {
    "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"
    }
  },
  "name": "lif1",
  "port_address": "5060F",
  "state": "up",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
```

```
"wvnn": "20:00:00:50:56:b4:13:01",  
"wvbn": "20:00:00:50:56:b4:13:a8"  
}
```

## 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	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

node

The node on which the FC port is located.

Name	Type	Description
name	string	The name of the node on which the FC port is located.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FC port.
node	<a href="#">node</a>	The node on which the FC port is located.
uuid	string	The unique identifier of the FC port.

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.

Name	Type	Description
node	<a href="#">node</a>	
port	<a href="#">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.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update an FC interface

PATCH /network/fc/interfaces/{uuid}

Updates an FC interface.

#### Related ONTAP commands

- `network interface modify`

#### Learn more

- [DOC /network/fc/interfaces](#)

#### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier for the FC interface.

#### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	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	The data protocol for which the FC interface is configured. Required in POST.
enabled	boolean	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 <i>true</i> (enabled) in POST.
location	<a href="#">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.

Name	Type	Description
name	string	The name of the FC interface. Required in POST; optional in PATCH.
port_address	string	<p>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).</p> <p>This is useful for obtaining statistics and diagnostic information from FC switches.</p> <p>This is a hexadecimal encoded numeric value.</p>
state	string	<p>The current operational state of the FC interface. The state is set to <i>down</i> if the interface is not enabled.</p> <p>If the node hosting the port is down or unavailable, no state value is returned.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the FC interface. Required in the URL.
wwnn	string	<p>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.</p> <ul style="list-style-type: none"> <li>example: 20:00:00:50:56:b4:13:01</li> <li>readOnly: 1</li> </ul>

Name	Type	Description
wwpn	string	<p>The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created.</p> <ul style="list-style-type: none"><li>• example: 20:00:00:50:56:b4:13:a8</li><li>• readOnly: 1</li></ul>

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "data_protocol": "fcp",
  "location": {
    "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"
    }
  },
  "name": "lif1",
  "port_address": "5060F",
  "state": "up",
  "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"  
}
```

## Response

Status: 200, Ok

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
1966140	A LIF 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 LIF 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 LIF 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 <code>location.port.node.name</code> or <code>location.port.name</code> is supplied, both properties must be supplied.
72089674	You cannot move a Fibre Channel interface configured for the NVMe over FC data protocol.

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	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

node

The node on which the FC port is located.

Name	Type	Description
name	string	The name of the node on which the FC port is located.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the FC port.
node	<a href="#">node</a>	The node on which the FC port is located.
uuid	string	The unique identifier of the FC port.

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.

Name	Type	Description
node	<a href="#">node</a>	
port	<a href="#">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.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	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	The data protocol for which the FC interface is configured. Required in POST.

Name	Type	Description
enabled	boolean	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 <i>true</i> (enabled) in POST.
location	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.
name	string	The name of the FC interface. Required in POST; optional in PATCH.
port_address	string	<p>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).</p> <p>This is useful for obtaining statistics and diagnostic information from FC switches.</p> <p>This is a hexadecimal encoded numeric value.</p>
state	string	<p>The current operational state of the FC interface. The state is set to <i>down</i> if the interface is not enabled.</p> <p>If the node hosting the port is down or unavailable, no state value is returned.</p>

Name	Type	Description
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the FC interface. Required in the URL.
wwnn	string	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. <ul style="list-style-type: none"> <li>example: 20:00:00:50:56:b4:13:01</li> <li>readOnly: 1</li> </ul>
wwpn	string	The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created. <ul style="list-style-type: none"> <li>example: 20:00:00:50:56:b4:13:a8</li> <li>readOnly: 1</li> </ul>

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve FC port information

### 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 are not supported. You must physically add and remove FC adapters to ONTAP nodes to create and remove ports from the ONTAP cluster.

#### 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"
        }
      }
    }
  },
  {
    "node": {
      "name": "0a",
      "state": "online",
      "uuid": "931b20f8-b047-11e8-9af3-005056bb838e",
      "_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"
        }
      }
    }
  },
  {
    "node": {
      "name": "0b",
      "state": "online",
      "uuid": "931b23f7-b047-11e8-9af3-005056bb838e",
      "_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": [
  "fcp"
],
"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",
"_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 }'
```

## Retrieve FC ports

GET /network/fc/ports

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 [DOC Requesting specific fields](#) to learn more.

- `fabric.name`

### Related ONTAP commands

- `network fcp adapter show`

### Learn more

- [DOC /network/fc/ports](#)

### Parameters

Name	Type	In	Required	Description
wwpn	string	query	False	Filter by wwpn
description	string	query	False	Filter by description
uuid	string	query	False	Filter by uuid

Name	Type	In	Required	Description
supported_protocols	string	query	False	Filter by supported_protocols
transceiver.manufacturer	string	query	False	Filter by transceiver.manufacturer
transceiver.form-factor	string	query	False	Filter by transceiver.form-factor
transceiver.part_number	string	query	False	Filter by transceiver.part_number
transceiver.capabilities	integer	query	False	Filter by transceiver.capabilities
wwnn	string	query	False	Filter by wwnn
speed.configured	string	query	False	Filter by speed.configured
speed.maximum	string	query	False	Filter by speed.maximum
enabled	boolean	query	False	Filter by enabled
fabric.connected	boolean	query	False	Filter by fabric.connected
fabric.switch_port	string	query	False	Filter by fabric.switch_port
fabric.connected_speed	integer	query	False	Filter by fabric.connected_speed
fabric.port_address	string	query	False	Filter by fabric.port_address
fabric.name	string	query	False	Filter by fabric.name
node.name	string	query	False	Filter by node.name

Name	Type	In	Required	Description
node.uuid	string	query	False	Filter by node.uuid
state	string	query	False	Filter by state
name	string	query	False	Filter by name
physical_protocol	string	query	False	Filter by physical_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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
num_records	integer	Number of records.
records	array[ <a href="#">fc_port</a> ]	

## Example response

```
{
  "_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"
    },
    "name": "0a",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "physical_protocol": "fibre_channel",
    "speed": {
      "configured": "auto",
      "maximum": 32
    },
    "state": "online",
    "supported_protocols": {
    },
    "transceiver": {
      "capabilities": {
    },
  },
}
```

```
    "form-factor": "sfp",
    "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

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

fabric

Properties of the fabric to which the FC port is attached.

Name	Type	Description
connected	boolean	Reports if the physical port has established a connection with the FC fabric.
connected_speed	integer	The negotiated data rate between the target FC port and the fabric in gigabits per second.
name	string	<p>The name of the fabric to which the port is connected. This is only available when the FC port is connected to a fabric.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
port_address	string	<p>The FC port address of the host bus adapter (HBA) physical port.</p> <p>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).</p> <p>This is useful for obtaining statistics and diagnostic information from FC switches.</p> <p>This is a six-digit hexadecimal encoded numeric value.</p>
switch_port	string	The switch port to which the FC port is connected.

#### node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

#### speed

The physical device speed related properties of the FC port.

Name	Type	Description
configured	string	The configured speed of the FC port in gigabits per second.
maximum	string	The maximum speed supported by the FC port in gigabits per second.

#### transceiver

Properties of the transceiver connected to the FC port.

Name	Type	Description
capabilities	array[integer]	The speeds of which the transceiver is capable in gigabits per second.
form-factor	string	The form factor of the transceiver. Possible values are: <ul style="list-style-type: none"> <li>• <i>sfp</i> - Small Form Factor - Pluggable</li> <li>• <i>sff</i> - Small Form Factor</li> <li>• <i>unk</i> - Unknown</li> </ul>
manufacturer	string	The manufacturer of the transceiver.
part_number	string	The part number of the transceiver.

#### 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.

Name	Type	Description
_links	<a href="#">_links</a>	
description	string	A description of the FC port.
enabled	boolean	The administrative state of the FC port. If this property is set to <i>false</i> , all FC connectivity to FC interfaces are blocked. Optional in PATCH.
fabric	<a href="#">fabric</a>	Properties of the fabric to which the FC port is attached.
name	string	The FC port name.
node	<a href="#">node</a>	
physical_protocol	string	The physical network protocol of the FC port.

Name	Type	Description
speed	speed	The physical device speed related properties of the FC port.
state	string	<p>The operational state of the FC port.</p> <ul style="list-style-type: none"> <li>• startup - The port is booting up.</li> <li>• link_not_connected - The port has finished initialization, but a link with the fabric is not established.</li> <li>• online - The port is initialized and a link with the fabric has been established.</li> <li>• link_disconnected - The link was present at one point on this port but is currently not established.</li> <li>• offlined_by_user - The port is administratively disabled.</li> <li>• offlined_by_system - The port is set to offline by the system. This happens when the port encounters too many errors.</li> <li>• node_offline - The state information for the port cannot be retrieved. The node is offline or inaccessible.</li> </ul>
supported_protocols	array[string]	The network protocols supported by the FC port.
transceiver	transceiver	Properties of the transceiver connected to the FC port.
uuid	string	The unique identifier of the FC port.
wwnn	string	The base world wide node name (WWNN) for the FC port.
wwpn	string	The base world wide port name (WWPN) for the FC port.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an FC port

GET /network/fc/ports/{uuid}

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 [DOC Requesting specific fields](#) to learn more.

- `fabric.name`

### Related ONTAP commands

- `network fcp adapter show`

### Learn more

- [DOC /network/fc/ports](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier for the FC port.

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

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
description	string	A description of the FC port.
enabled	boolean	The administrative state of the FC port. If this property is set to <i>false</i> , all FC connectivity to FC interfaces are blocked. Optional in PATCH.
<a href="#">fabric</a>	<a href="#">fabric</a>	Properties of the fabric to which the FC port is attached.
name	string	The FC port name.
<a href="#">node</a>	<a href="#">node</a>	
physical_protocol	string	The physical network protocol of the FC port.
speed	<a href="#">speed</a>	The physical device speed related properties of the FC port.

Name	Type	Description
state	string	<p>The operational state of the FC port.</p> <ul style="list-style-type: none"> <li>• startup - The port is booting up.</li> <li>• link_not_connected - The port has finished initialization, but a link with the fabric is not established.</li> <li>• online - The port is initialized and a link with the fabric has been established.</li> <li>• link_disconnected - The link was present at one point on this port but is currently not established.</li> <li>• offlined_by_user - The port is administratively disabled.</li> <li>• offlined_by_system - The port is set to offline by the system. This happens when the port encounters too many errors.</li> <li>• node_offline - The state information for the port cannot be retrieved. The node is offline or inaccessible.</li> </ul>
supported_protocols	array[string]	The network protocols supported by the FC port.
transceiver	<a href="#">transceiver</a>	Properties of the transceiver connected to the FC port.
uuid	string	The unique identifier of the FC port.
wwnn	string	The base world wide node name (WWNN) for the FC port.
wwpn	string	The base world wide port name (WWPN) for the FC port.

## Example response

```
{
  "_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"
  },
  "name": "0a",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "physical_protocol": "fibre_channel",
  "speed": {
    "configured": "auto",
    "maximum": 32
  },
  "state": "online",
  "supported_protocols": {
  },
  "transceiver": {
    "capabilities": {
    },
    "form-factor": "sfp",
    "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

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	<a href="#">href</a>	

fabric

Properties of the fabric to which the FC port is attached.

Name	Type	Description
connected	boolean	Reports if the physical port has established a connection with the FC fabric.
connected_speed	integer	The negotiated data rate between the target FC port and the fabric in gigabits per second.
name	string	<p>The name of the fabric to which the port is connected. This is only available when the FC port is connected to a fabric.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
port_address	string	<p>The FC port address of the host bus adapter (HBA) physical port.</p> <p>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).</p> <p>This is useful for obtaining statistics and diagnostic information from FC switches.</p> <p>This is a six-digit hexadecimal encoded numeric value.</p>
switch_port	string	The switch port to which the FC port is connected.

#### node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

#### speed

The physical device speed related properties of the FC port.

Name	Type	Description
configured	string	The configured speed of the FC port in gigabits per second.
maximum	string	The maximum speed supported by the FC port in gigabits per second.

#### transceiver

Properties of the transceiver connected to the FC port.

Name	Type	Description
capabilities	array[integer]	The speeds of which the transceiver is capable in gigabits per second.
form-factor	string	The form factor of the transceiver. Possible values are: <ul style="list-style-type: none"> <li>• <i>sfp</i> - Small Form Factor - Pluggable</li> <li>• <i>sff</i> - Small Form Factor</li> <li>• <i>unk</i> - Unknown</li> </ul>
manufacturer	string	The manufacturer of the transceiver.
part_number	string	The part number of the transceiver.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update an FC port

PATCH /network/fc/ports/{uuid}

Updates an FC port.

## Related ONTAP commands

- `network fcp adapter modify`

## Learn more

- [DOC /network/fc/ports](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier for the FC port.

## Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
description	string	A description of the FC port.
enabled	boolean	The administrative state of the FC port. If this property is set to <i>false</i> , all FC connectivity to FC interfaces are blocked. Optional in PATCH.
fabric	<a href="#">fabric</a>	Properties of the fabric to which the FC port is attached.
name	string	The FC port name.
node	<a href="#">node</a>	
physical_protocol	string	The physical network protocol of the FC port.
speed	<a href="#">speed</a>	The physical device speed related properties of the FC port.

Name	Type	Description
state	string	<p>The operational state of the FC port.</p> <ul style="list-style-type: none"> <li>• startup - The port is booting up.</li> <li>• link_not_connected - The port has finished initialization, but a link with the fabric is not established.</li> <li>• online - The port is initialized and a link with the fabric has been established.</li> <li>• link_disconnected - The link was present at one point on this port but is currently not established.</li> <li>• offlined_by_user - The port is administratively disabled.</li> <li>• offlined_by_system - The port is set to offline by the system. This happens when the port encounters too many errors.</li> <li>• node_offline - The state information for the port cannot be retrieved. The node is offline or inaccessible.</li> </ul>
supported_protocols	array[string]	The network protocols supported by the FC port.
transceiver	<a href="#">transceiver</a>	Properties of the transceiver connected to the FC port.
uuid	string	The unique identifier of the FC port.
wwnn	string	The base world wide node name (WWNN) for the FC port.
wwpn	string	The base world wide port name (WWPN) for the FC port.

## Example request

```
{
  "_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"
  },
  "name": "0a",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "physical_protocol": "fibre_channel",
  "speed": {
    "configured": "auto",
    "maximum": 32
  },
  "state": "online",
  "supported_protocols": {
  },
  "transceiver": {
    "capabilities": {
    },
    "form-factor": "sfp",
    "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.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

fabric

Properties of the fabric to which the FC port is attached.

Name	Type	Description
connected	boolean	Reports if the physical port has established a connection with the FC fabric.
connected_speed	integer	The negotiated data rate between the target FC port and the fabric in gigabits per second.
name	string	<p>The name of the fabric to which the port is connected. This is only available when the FC port is connected to a fabric.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
port_address	string	<p>The FC port address of the host bus adapter (HBA) physical port.</p> <p>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).</p> <p>This is useful for obtaining statistics and diagnostic information from FC switches.</p> <p>This is a six-digit hexadecimal encoded numeric value.</p>
switch_port	string	The switch port to which the FC port is connected.

#### node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

#### speed

The physical device speed related properties of the FC port.

Name	Type	Description
configured	string	The configured speed of the FC port in gigabits per second.
maximum	string	The maximum speed supported by the FC port in gigabits per second.

#### transceiver

Properties of the transceiver connected to the FC port.

Name	Type	Description
capabilities	array[integer]	The speeds of which the transceiver is capable in gigabits per second.
form-factor	string	The form factor of the transceiver. Possible values are: <ul style="list-style-type: none"> <li>• <i>sfp</i> - Small Form Factor - Pluggable</li> <li>• <i>sff</i> - Small Form Factor</li> <li>• <i>unk</i> - Unknown</li> </ul>
manufacturer	string	The manufacturer of the transceiver.
part_number	string	The part number of the transceiver.

#### 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.

Name	Type	Description
_links	<a href="#">_links</a>	
description	string	A description of the FC port.
enabled	boolean	The administrative state of the FC port. If this property is set to <i>false</i> , all FC connectivity to FC interfaces are blocked. Optional in PATCH.
fabric	<a href="#">fabric</a>	Properties of the fabric to which the FC port is attached.
name	string	The FC port name.
node	<a href="#">node</a>	
physical_protocol	string	The physical network protocol of the FC port.

Name	Type	Description
speed	speed	The physical device speed related properties of the FC port.
state	string	<p>The operational state of the FC port.</p> <ul style="list-style-type: none"> <li>• startup - The port is booting up.</li> <li>• link_not_connected - The port has finished initialization, but a link with the fabric is not established.</li> <li>• online - The port is initialized and a link with the fabric has been established.</li> <li>• link_disconnected - The link was present at one point on this port but is currently not established.</li> <li>• offlined_by_user - The port is administratively disabled.</li> <li>• offlined_by_system - The port is set to offline by the system. This happens when the port encounters too many errors.</li> <li>• node_offline - The state information for the port cannot be retrieved. The node is offline or inaccessible.</li> </ul>
supported_protocols	array[string]	The network protocols supported by the FC port.
transceiver	transceiver	Properties of the transceiver connected to the FC port.
uuid	string	The unique identifier of the FC port.
wwnn	string	The base world wide node name (WWNN) for the FC port.
wwpn	string	The base world wide port name (WWPN) for the FC port.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage network IP interfaces

### 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",
"ipSPACE": {
  "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.ho
me_node.name,service_policy.name,ip.address,enabled"
  }
}
}
}

```

## Creating IP interfaces

The IP interfaces POST API is used 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

The IP interfaces PATCH API is used to update attributes of 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": "default-management" }'
```

---

## Deleting IP interfaces

The IP interfaces DELETE API is used 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"
```

---

## Retrieve all IP interface details

GET /network/ip/interfaces

Retrieves the details of all IP interfaces.

## Related ONTAP Commands

- `network interface show`

## Learn more

- [DOC /network/ip/interfaces](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
scope	string	query	False	Filter by scope
services	string	query	False	Filter by services
state	string	query	False	Filter by state
name	string	query	False	Filter by name
enabled	boolean	query	False	Filter by enabled
ipspace.uuid	string	query	False	Filter by ipspace.uuid
ipspace.name	string	query	False	Filter by ipspace.name
ip.address	string	query	False	Filter by ip.address
ip.netmask	string	query	False	Filter by ip.netmask
ip.family	string	query	False	Filter by ip.family
vip	boolean	query	False	Filter by vip
service_policy.uuid	string	query	False	Filter by service_policy.uuid
service_policy.name	string	query	False	Filter by service_policy.name
uuid	string	query	False	Filter by uuid

<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
location.home_node.name	string	query	False	Filter by location.home_node.name
location.home_node.uuid	string	query	False	Filter by location.home_node.uuid
location.is_home	boolean	query	False	Filter by location.is_home
location.auto_revert	boolean	query	False	Filter by location.auto_revert
location.failover	string	query	False	Filter by location.failover
location.home_port.node.name	string	query	False	Filter by location.home_port.node.name
location.home_port.name	string	query	False	Filter by location.home_port.name
location.home_port.uuid	string	query	False	Filter by location.home_port.uuid
location.node.name	string	query	False	Filter by location.node.name
location.node.uuid	string	query	False	Filter by location.node.uuid
location.port.node.name	string	query	False	Filter by location.port.node.name
location.port.name	string	query	False	Filter by location.port.name
location.port.uuid	string	query	False	Filter by location.port.uuid

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	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	
records	array[ <a href="#">ip_interface</a> ]	

## Example response

```
{
  "_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": "ipv4",
      "netmask": "24"
    },
    "ipspace": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "exchange",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "location": {
      "broadcast_domain": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "bd1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "failover": "home_port_only",
      "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"
}
},
"name": "dataLif1",
"scope": "svm",
"service_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
}
```

```

    }
  },
  "name": "default-intercluster",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"services": {
},
"state": "up",
"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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

ip\_info

IP information

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.

ipspace

Either the UUID or name must be supplied on POST for cluster-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

broadcast\_domain

Broadcast domain UUID along with a readable name.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	Name of the broadcast domain, scoped to its IPspace
uuid	string	Broadcast domain UUID

home\_node

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

node

Name	Type	Description
name	string	Name of node on which the port is located.

home\_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

node

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

#### 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.

Name	Type	Description
auto_revert	boolean	
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name.
failover	string	Defines where an interface may failover.
home_node	<a href="#">home_node</a>	
home_port	<a href="#">home_port</a>	Port UUID along with readable names. Either the UUID or both names may be supplied on input.
is_home	boolean	
node	<a href="#">node</a>	
port	<a href="#">port</a>	Port UUID along with readable names. Either the UUID or both names may be supplied on input.

#### service\_policy

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

#### svm

Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### ip\_interface

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	The administrative state of the interface.
ip	<a href="#">ip_info</a>	IP information
ipspace	<a href="#">ipspace</a>	Either the UUID or name must be supplied on POST for cluster-scoped objects.
location	<a href="#">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	Interface name
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
service_policy	<a href="#">service_policy</a>	
services	array[string]	The services associated with the interface.
state	string	The operational state of the interface.

Name	Type	Description
svm	<a href="#">svm</a>	Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.
uuid	string	The UUID that uniquely identifies the interface.
vip	boolean	True for a VIP interface, whose location is announced via BGP.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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 interface

POST /network/ip/interfaces

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 a 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.

### 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`

### Learn more

- [DOC /network/ip/interfaces](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>enabled</code>	boolean	The administrative state of the interface.
<code>ip</code>	<a href="#">ip_info</a>	IP information
<code>ipspace</code>	<a href="#">ipspace</a>	Either the UUID or name must be supplied on POST for cluster-scoped objects.

Name	Type	Description
location	<a href="#">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	Interface name
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
service_policy	<a href="#">service_policy</a>	
services	array[string]	The services associated with the interface.
state	string	The operational state of the interface.
svm	<a href="#">svm</a>	Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.
uuid	string	The UUID that uniquely identifies the interface.
vip	boolean	True for a VIP interface, whose location is announced via BGP.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip": {
    "address": "10.10.10.7",
    "family": "ipv4",
    "netmask": "24"
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "location": {
    "broadcast_domain": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "bd1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "failover": "home_port_only",
    "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"
}
},
"name": "dataLif1",
"scope": "svm",
"service_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default-intercluster",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"services": {
},
"state": "up",
"svm": {
  "_links": {

```

```

    "self": {
      "href": "/api/resourcelink"
    },
    "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
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".
1967082	The specified ipspace.name does not match the IPspace name of ipspace.uuid.
1967081	The specified SVM must exist in the specified IPspace.
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	Invalid location.broadcast_domain.uuid specified.
1967106	The specified location.home_port.name does not match the specified port name of location.home_port.uuid.

Error Code	Description
1967107	Invalid location.home_port.uuid specified.
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.
1967102	POST operation might have left configuration in an inconsistent state. Check the configuration.
1967137	The specified location.broadcast_domain.uuid and location.broadcast_domain.name (and IPspace name) are invalid.
1967135	The specified location.broadcast_domain.uuid is invalid.
1967136	The specified location.broadcast_domain.name (and ipspace name) is invalid.
1967129	The specified location.home_port.uuid is invalid.
1967130	The specified location.home_port.name is invalid.
1967131	The specified location.home_port.uuid and location.home_port.name are invalid.
1967145	The specified location.failover is invalid.
1967146	The specified svm.name is invalid.
1967147	The specified svm.uuid is invalid.
5373966	An iSCSI interface cannot be created in an SVM is configured for NVMe.
1966140	A LIF with the same name already exists.
1966987	The Vserver Broadcast-Domain Home-Node and Home-Port combination is not valid.
1376656	Cluster LIFs must be in the same subnet. Verify the address and netmask are set to the correct values.
1966138	The same IP address may not be used for both a mgmt interface and a gateway address.
1376963	Duplicate IP address

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	<a href="#">href</a>	

ip\_info

IP information

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.

ipspace

Either the UUID or name must be supplied on POST for cluster-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

broadcast\_domain

Broadcast domain UUID along with a readable name.

Name	Type	Description
_links	<a href="#">_links</a>	



Name	Type	Description
name	string	Name of the broadcast domain, scoped to its IPspace
uuid	string	Broadcast domain UUID

home\_node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

node

Name	Type	Description
name	string	Name of node on which the port is located.

home\_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	
node	<a href="#">node</a>	
uuid	string	

#### 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.

Name	Type	Description
auto_revert	boolean	
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name.
failover	string	Defines where an interface may failover.
home_node	<a href="#">home_node</a>	
home_port	<a href="#">home_port</a>	Port UUID along with readable names. Either the UUID or both names may be supplied on input.
is_home	boolean	
node	<a href="#">node</a>	
port	<a href="#">port</a>	Port UUID along with readable names. Either the UUID or both names may be supplied on input.

#### service\_policy

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

#### svm

Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### ip\_interface

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	The administrative state of the interface.
ip	<a href="#">ip_info</a>	IP information
ipspace	<a href="#">ipspace</a>	Either the UUID or name must be supplied on POST for cluster-scoped objects.
location	<a href="#">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	Interface name
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
service_policy	<a href="#">service_policy</a>	
services	array[string]	The services associated with the interface.
state	string	The operational state of the interface.
svm	<a href="#">svm</a>	Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

Name	Type	Description
uuid	string	The UUID that uniquely identifies the interface.
vip	boolean	True for a VIP interface, whose location is announced via BGP.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete an IP interface

DELETE /network/ip/interfaces/{uuid}

Deletes an IP interface.

### Related ONTAP commands

- `network interface delete`

### Learn more

- [DOC /network/ip/interfaces](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

## Response

```
Status: 200, Ok
```

## Retrieve details for an IP interface

```
GET /network/ip/interfaces/{uuid}
```

Retrieves details for a specific IP interface.

## Related ONTAP commands

- `network interface show`

## Learn more

- [DOC /network/ip/interfaces](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	IP interface 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>enabled</code>	boolean	The administrative state of the interface.
<code>ip</code>	<code>ip_info</code>	IP information
<code>ipspace</code>	<code>ipspace</code>	Either the UUID or name must be supplied on POST for cluster-scoped objects.

Name	Type	Description
location	<a href="#">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	Interface name
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
service_policy	<a href="#">service_policy</a>	
services	array[string]	The services associated with the interface.
state	string	The operational state of the interface.
svm	<a href="#">svm</a>	Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.
uuid	string	The UUID that uniquely identifies the interface.
vip	boolean	True for a VIP interface, whose location is announced via BGP.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip": {
    "address": "10.10.10.7",
    "family": "ipv4",
    "netmask": "24"
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "location": {
    "broadcast_domain": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "bd1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "failover": "home_port_only",
    "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"
}
},
"name": "dataLif1",
"scope": "svm",
"service_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default-intercluster",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"services": {
},
"state": "up",
"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	<a href="#">error</a>	

## 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	<a href="#">href</a>	

ip\_info

IP information

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.

ipspace

Either the UUID or name must be supplied on POST for cluster-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

broadcast\_domain

Broadcast domain UUID along with a readable name.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	Name of the broadcast domain, scoped to its IPspace
uuid	string	Broadcast domain UUID

home\_node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

node

Name	Type	Description
name	string	Name of node on which the port is located.

home\_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	
node	<a href="#">node</a>	
uuid	string	

#### 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.

Name	Type	Description
auto_revert	boolean	
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name.
failover	string	Defines where an interface may failover.
home_node	<a href="#">home_node</a>	
home_port	<a href="#">home_port</a>	Port UUID along with readable names. Either the UUID or both names may be supplied on input.
is_home	boolean	
node	<a href="#">node</a>	
port	<a href="#">port</a>	Port UUID along with readable names. Either the UUID or both names may be supplied on input.

#### service\_policy

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

#### svm

Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

Name	Type	Description
_links	<a href="#">_links</a>	

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update an IP interface

PATCH /network/ip/interfaces/{uuid}

Updates an IP interface.

### Related ONTAP commands

- `network interface migrate`
- `network interface modify`
- `network interface rename`
- `network interface revert`

### Learn more

- [DOC /network/ip/interfaces](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	IP interface UUID

#### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	The administrative state of the interface.
ip	<a href="#">ip_info</a>	IP information
ipspace	<a href="#">ipspace</a>	Either the UUID or name must be supplied on POST for cluster-scoped objects.
location	<a href="#">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	Interface name
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
service_policy	<a href="#">service_policy</a>	
services	array[string]	The services associated with the interface.
state	string	The operational state of the interface.
svm	<a href="#">svm</a>	Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.
uuid	string	The UUID that uniquely identifies the interface.

<b>Name</b>	<b>Type</b>	<b>Description</b>
vip	boolean	True for a VIP interface, whose location is announced via BGP.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip": {
    "address": "10.10.10.7",
    "family": "ipv4",
    "netmask": "24"
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "location": {
    "broadcast_domain": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "bd1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "failover": "home_port_only",
    "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"
}
},
"name": "dataLif1",
"scope": "svm",
"service_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default-intercluster",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"services": {
},
"state": "up",
"svm": {
  "_links": {

```

```

    "self": {
      "href": "/api/resourcelink"
    },
    "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
1967143	Cannot patch home_port for a VIP interface. Invalid parameters location.home_port.uuid and location.home_port.name specified.
1967141	Cannot patch home_port for a VIP interface. Invalid parameter location.home_port.uuid specified.
1967142	Cannot patch home_port for a VIP interface. Invalid parameter location.home_port.name specified.
1967140	Cannot patch port for a VIP interface. Invalid parameters location.port.uuid and location.port.name specified.
1967138	Cannot patch port for a VIP interface. Invalid parameter location.port.uuid specified.
1967139	Cannot patch port for a VIP interface. Invalid parameter location.port.name specified.
1967120	The specified service_policy.name does not match the specified service policy name of service_policy.uuid.
1967121	Invalid service_policy.uuid specified.
1967113	The specified location.port.name does not match the port name of location.port.uuid.
1967114	Invalid location.port.uuid specified.

Error Code	Description
1967134	The specified location.port.uuid and location.port.name are invalid.
1967132	The specified location.port.uuid is invalid.
1967133	The specified location.port.name is invalid.
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	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.
1967106	The specified location.home_port.name does not match the specified port name of location.home_port.uuid.
1967107	Invalid location.home_port.uuid specified.
1967111	Home node must be specified by at least one location.home_node, location.home_port, or location.broadcast_domain field.
1967129	The specified location.home_port.uuid is invalid.
1967130	The specified location.home_port.name is invalid.
1967131	The specified location.home_port.uuid and location.home_port.name are invalid.
1967145	The specified location.failover is invalid.
1966138	The same IP address may not be used for both a mgmt interface and a gateway address.
1376963	Duplicate IP address
1966229	You cannot patch the "location.node" or "location.port" fields to migrate interfaces using FCP and iSCSI data protocols. 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".

## Definitions

## See Definitions

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

ip\_info

IP information

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.

ipspace

Either the UUID or name must be supplied on POST for cluster-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

broadcast\_domain

Broadcast domain UUID along with a readable name.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	Name of the broadcast domain, scoped to its IPspace
uuid	string	Broadcast domain UUID

home\_node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

node

Name	Type	Description
name	string	Name of node on which the port is located.

home\_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	
uuid	string	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	
node	<a href="#">node</a>	
uuid	string	

#### 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.

Name	Type	Description
auto_revert	boolean	
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name.
failover	string	Defines where an interface may failover.
home_node	<a href="#">home_node</a>	
home_port	<a href="#">home_port</a>	Port UUID along with readable names. Either the UUID or both names may be supplied on input.
is_home	boolean	
node	<a href="#">node</a>	
port	<a href="#">port</a>	Port UUID along with readable names. Either the UUID or both names may be supplied on input.

#### service\_policy

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

#### svm

Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### ip\_interface

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	The administrative state of the interface.
ip	<a href="#">ip_info</a>	IP information
ipspace	<a href="#">ipspace</a>	Either the UUID or name must be supplied on POST for cluster-scoped objects.
location	<a href="#">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	Interface name
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
service_policy	<a href="#">service_policy</a>	
services	array[string]	The services associated with the interface.
state	string	The operational state of the interface.
svm	<a href="#">svm</a>	Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

Name	Type	Description
uuid	string	The UUID that uniquely identifies the interface.
vip	boolean	True for a VIP interface, whose location is announced via BGP.

## Manage network IP routes

### Network IP routes endpoint overview

#### Overview

This endpoint supports the following operations: GET (collection and instance), POST, and DELETE.

#### Retrieving network routes

The IP routes GET API retrieves and displays 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",
      "ipospace": {
        "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

# 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",
  "ipospace": {
    "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

The POST API is used 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

The DELETE API can remove a specific route identified by its UUID.

#### Example

#### Deleting a specific route

```
# The API:
/api/network/ip/routes

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/ip/routes/38805a91-9ac9-
11e8-8bc9-005056bbd531"
```

---

### Retrieve IP routes

GET /network/ip/routes

Retrieves the collection of IP routes.

#### Related ONTAP commands

- `network route show`

#### Learn more

- [DOC /network/ip/routes](#)

#### Parameters

Name	Type	In	Required	Description
ipspace.uuid	string	query	False	Filter by ipspace.uuid
ipspace.name	string	query	False	Filter by ipspace.name
uuid	string	query	False	Filter by uuid
destination.address	string	query	False	Filter by destination.address
destination.netmask	string	query	False	Filter by destination.netmask
destination.family	string	query	False	Filter by destination.family
gateway	string	query	False	Filter by gateway
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
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_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
error	<a href="#">error</a>	
num_records	integer	Number of records
records	array[ <a href="#">network_route</a> ]	

## 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": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "destination": {
      "address": "10.10.10.7",
      "family": "ipv4",
      "netmask": "24"
    },
    "gateway": "10.1.1.1",
    "ipspace": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "exchange",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "scope": "svm",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

\_links

Name	Type	Description
self	<a href="#">href</a>	

ip\_info

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

Name	Type	Description
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.

#### ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### network\_route

Name	Type	Description
_links	<a href="#">_links</a>	
destination	<a href="#">ip_info</a>	IP information
gateway	string	The IP address of the gateway router leading to the destination.
ipspace	<a href="#">ipspace</a>	Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.

Name	Type	Description
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	The UUID that uniquely identifies the route.

### Create a cluster-scoped or SVM-scoped static route

POST /network/ip/routes

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

#### Learn more

- [DOC /network/ip/routes](#)

#### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
destination	ip_info	IP information
gateway	string	The IP address of the gateway router leading to the destination.
ipspace	ipspace	Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.
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	The UUID that uniquely identifies the route.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "destination": {
    "address": "10.10.10.7",
    "family": "ipv4",
    "netmask": "24"
  },
  "gateway": "10.1.1.1",
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "exchange",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"scope": "svm",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

## Response

Status: 201, Created

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	



Name	Type	Description
error	<a href="#">error</a>	
num_records	integer	Number of records
records	array[ <a href="#">network_route</a> ]	

## 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": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "destination": {
      "address": "10.10.10.7",
      "family": "ipv4",
      "netmask": "24"
    },
    "gateway": "10.1.1.1",
    "ipspace": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "exchange",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "scope": "svm",
    "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
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 invalid.
2	The specified svm.uuid is invalid or wrong.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

ip\_info

IP information

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.

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.

Name	Type	Description
uuid	string	The unique identifier of the SVM.

#### network\_route

Name	Type	Description
_links	<a href="#">_links</a>	
destination	<a href="#">ip_info</a>	IP information
gateway	string	The IP address of the gateway router leading to the destination.
ipspace	<a href="#">ipspace</a>	Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The UUID that uniquely identifies the route.

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete an IP route

DELETE /network/ip/routes/{uuid}

Deletes a specific IP route.

#### Related ONTAP commands

- `network route delete`

#### Learn more

- [DOC /network/ip/routes](#)

#### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

#### Response

Status: 200, Ok

#### Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve details for an IP route

GET /network/ip/routes/{uuid}

Retrieves the details of a specific IP route.



## Related ONTAP commands

- `network route show`

## Learn more

- [DOC /network/ip/routes](#)

## Parameters

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

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
destination	<a href="#">ip_info</a>	IP information
gateway	string	The IP address of the gateway router leading to the destination.
ipspace	<a href="#">ipspace</a>	Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The UUID that uniquely identifies the route.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "destination": {
    "address": "10.10.10.7",
    "family": "ipv4",
    "netmask": "24"
  },
  "gateway": "10.1.1.1",
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "exchange",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"scope": "svm",
"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	<a href="#">href</a>	

ip\_info

IP information

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.

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage network IP service policies

### Network IP service-policies endpoint overview

#### Overview

Service policies are named groupings that define what services are supported by an IP interface. The network IP service-policies GET API retrieves and displays relevant information pertaining to the service policies configured in the cluster. The API retrieves the list of all service policies configured in the cluster or a specific service policy.

#### 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

# 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"
      }
    }
  },
  "ipospace": {
    "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

# 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

# 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

# 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"
    }
  }
}
```

---

## Retrieve service policies

GET /network/ip/service-policies

Retrieves a collection of service policies.

### Related ONTAP commands

- `network interface service-policy show`

### Learn more

- [DOC /network/ip/service-policies](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	query	False	Filter by uuid
ipspace.uuid	string	query	False	Filter by ipspace.uuid
ipspace.name	string	query	False	Filter by ipspace.name
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
name	string	query	False	Filter by name
services	string	query	False	Filter by services
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	
records	array[ <a href="#">ip_service_policy</a> ]	

## Example response

```
{
  "_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": "svm",
    "services": {
    },
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ip\_service\_policy

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	



Name	Type	Description
ipspace	<a href="#">ipspace</a>	Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.
name	string	
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
services	array[string]	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Retrieve a service policy

GET /network/ip/service-policies/{uuid}

Retrieves a specific service policy.

## Related ONTAP commands

- `network interface service-policy show`

## Learn more

- [DOC /network/ip/service-policies](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Service policy 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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

```
Status: 200, Ok
```

Name	Type	Description
_links	<a href="#">_links</a>	
ipospace	<a href="#">ipospace</a>	Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.
name	string	
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
services	array[string]	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ipospace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "default-intercluster",
  "scope": "svm",
  "services": {
  },
  "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	<a href="#">href</a>	

ipospace

Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage network IPspaces

### 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

The IPspaces GET API retrieves 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

The network IPspaces POST API is used 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

The IPspaces PATCH API is used to modify 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

The IPspaces DELETE API is used 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"
```

---

## Retrieve IPspaces for a cluster

GET /network/ipspaces

Retrieves a collection of IPspaces for the entire cluster.

### Related ONTAP commands

- `network ipspace show`

### Learn more

- [DOC /network/ipspaces](#)

## Parameters

Name	Type	In	Required	Description
name	string	query	False	Filter by 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_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	
records	array[ <a href="#">ipospace</a> ]	

## Example response

```
{
  "_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

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

ipospace

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	The UUID that uniquely identifies the IPspace.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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 domain with unique IP addresses

POST /network/ipspaces

Creates a new domain within which IP addresses are unique. SVMs, ports, and networks are scoped within a single IPspace.

### Required properties

- name - Name of the ipspace to create.

### Related ONTAP commands

- `network ipspace create`

### Learn more

- [DOC /network/ipspaces](#)

### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	The UUID that uniquely identifies the IPspace.

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resource/link"
    }
  },
  "name": "ipspace1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```



## Response

Status: 201, Created

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
1967102	POST operation might have left the configuration in an inconsistent state. Check the configuration.

## ONTAP Error Response Codes

Error Code	Description
9240591	Invalid name. The name is already in use by a cluster node, Vserver, or is the name of the local 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	<a href="#">href</a>	

ipospace

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	The UUID that uniquely identifies the IPspace.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete an IPspace object

DELETE /network/ipspaces/{uuid}

Deletes an IPspace object.

#### Related ONTAP commands

- `network ipspace delete`

#### Learn more

- [DOC /network/ipspaces](#)

#### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

#### Response

```
Status: 200, Ok
```

#### Retrieve information about an IPspace

GET /network/ipspaces/{uuid}

Retrieves information about a specific IPspace.

#### Related ONTAP commands

- `network ipspace show`

#### Learn more

- [DOC /network/ipspaces](#)

#### Parameters

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

#### Response

```
Status: 200, Ok
```

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	The UUID that uniquely identifies the IPspace.

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "ipspace1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

### Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

### 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	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Update an IPspace object

```
PATCH /network/ipspaces/{uuid}
```

Updates an IPspace object.

#### Related ONTAP commands

- `network ipspace rename`

#### Learn more

- [DOC /network/ipspaces](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	IPspace UUID

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	The UUID that uniquely identifies the IPspace.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resource/link"
    }
  },
  "name": "ipspace1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

## Response

```
Status: 200, Ok
```

## Definitions

## See Definitions

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

ipspace

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	The UUID that uniquely identifies the IPspace.

# SAN

## 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

### 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": {
```



```

    },
    "_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"
  }
}
}
}

```

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=fc" -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": "fc",
      "_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": "svml",
    "_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": "svml",
      "_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": "fc",
  "_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=*"
  }
}
}
}

```

## Retrieve FC logins

GET /network/fc/logins

Retrieves FC logins.

### Related ONTAP commands

- `vserver fcp initiator show`

### Learn more

- SAN: [DOC /network/fc/logins](#)
- NVMe: [DOC /network/fc/logins](#)

### Parameters

Name	Type	In	Required	Description
igroups.name	string	query	False	Filter by igroups.name
igroups.uuid	string	query	False	Filter by igroups.uuid
initiator.port_address	string	query	False	Filter by initiator.port_addresses

<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
initiator.wwnn	string	query	False	Filter by initiator.wwnn
initiator.wwpn	string	query	False	Filter by initiator.wwpn
initiator.aliases	string	query	False	Filter by initiator.aliases
protocol	string	query	False	Filter by protocol
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
interface.name	string	query	False	Filter by interface.name
interface.wwpn	string	query	False	Filter by interface.wwpn
interface.uuid	string	query	False	Filter by interface.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.



Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">fc_login</a> ]	

## Example response

```
{
  "_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": {
      },
      "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": "fc_nvme",
"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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

igroups

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the initiator group.
uuid	string	The unique identifier of the initiator group.

initiator

Information about the logged in FC initiator.

Name	Type	Description
aliases	array[string]	The logged in initiator world wide port name (WWPN) aliases.

Name	Type	Description
port_address	string	<p>The port address of the initiator's FC port.</p> <p>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.</p> <p>This is useful for obtaining statistics and diagnostic information from FC switches.</p> <p>This is a hexadecimal encoded numeric value.</p>
wwnn	string	The logged in initiator world wide node name (WWNN).
wwpn	string	The logged in initiator WWPN.

#### interface

An FC interface.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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.

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.

Name	Type	Description
uuid	string	The unique identifier of the SVM.

### 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.

Name	Type	Description
_links	<a href="#">_links</a>	
igroups	array[ <a href="#">igroups</a> ]	The initiator groups in which the initiator is a member.
initiator	<a href="#">initiator</a>	Information about the logged in FC initiator.
interface	<a href="#">interface</a>	An FC interface.
protocol	string	The data protocol used to perform the login.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an FC login

GET /network/fc/logins/{interface.uuid}/{initiator.wwpn}

Retrieves an FC login.

### Related ONTAP commands

- `vserver fcp initiator show`

### Learn more

- SAN: [DOC /network/fc/logins](#)
- NVMe: [DOC /network/fc/logins](#)

### Parameters

Name	Type	In	Required	Description
interface.uuid	string	path	True	The unique identifier of the FC interface through which the initiator logged in.
initiator.wwpn	string	path	True	The world wide port name (WWPN) of the initiator.
fields	array[string]	query	False	Specify the fields to return.

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
igroups	array[ <a href="#">igroups</a> ]	The initiator groups in which the initiator is a member.
initiator	<a href="#">initiator</a>	Information about the logged in FC initiator.
interface	<a href="#">interface</a>	An FC interface.
protocol	string	The data protocol used to perform the login.

Name	Type	Description
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.



## Example response

```
{
  "_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": {
    },
    "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": "fc_nvme",
  "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.
5374881	The Fibre Channel interface specified does not exist.
5373983	An invalid WWPN was supplied.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

igroups

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the initiator group.
uuid	string	The unique identifier of the initiator group.

initiator

Information about the logged in FC initiator.

Name	Type	Description
aliases	array[string]	The logged in initiator world wide port name (WWPN) aliases.

Name	Type	Description
port_address	string	<p>The port address of the initiator's FC port.</p> <p>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.</p> <p>This is useful for obtaining statistics and diagnostic information from FC switches.</p> <p>This is a hexadecimal encoded numeric value.</p>
wwnn	string	The logged in initiator world wide node name (WWNN).
wwpn	string	The logged in initiator WWPN.

#### interface

An FC interface.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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.

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage FC WWPN aliases

### Network FC wwpn-aliases endpoint overview

#### Overview

A world wide 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 may 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 of the properties of all the 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 with the alias "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"
```

## Retrieve FC WWPN aliases

GET /network/fc/wwpn-aliases

Retrieves FC WWPN aliases.

### Related ONTAP commands

- `vserver fcp wwpn-alias show`

### Learn more

- [DOC /network/fc/wwpn-aliases](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
alias	string	query	False	Filter by alias
wwpn	string	query	False	Filter by wwpn
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.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">wwpn_alias</a> ]	

## Example response

```
{
  "_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

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
alias	string	The FC WWPN alias. Required in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
wwpn	string	The FC initiator WWPN. Required in POST.

## error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

## error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create an FC WWPN alias

POST `/network/fc/wwpn-aliases`

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

- [DOC /network/fc/wwpn-aliases](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
alias	string	The FC WWPN alias. Required in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
wwpn	string	The FC initiator WWPN. Required in POST.

### Example request

```
{
  "_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"
}
```

### Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">wwpn_alias</a> ]	



## Example response

```
{
  "_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

## 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.

Error Code	Description
2621706	Both the SVM UUID and SVM name were supplied, but don't refer to the same SVM.
2621707	No SVM was specified. Either <code>svm.name</code> or <code>svm.uuid</code> must be supplied.
5373982	An invalid WWPN was supplied. The valid WWN format is <code>XX:XX:XX:XX:XX:XX:XX</code> , where X is a hexadecimal digit. Example: <code>"01:02:03:04:0a:0b:0c:0d"</code> .

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
alias	string	The FC WWPN alias. Required in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
wwpn	string	The FC initiator WWPN. Required in POST.

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

## error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

## error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete an FC WWPN alias

```
DELETE /network/fc/wwpn-aliases/{svm.uuid}/{alias}
```

Deletes an FC WWPN alias.

### Related ONTAP commands

- `vserver fcp wwpn-alias remove`

### Learn more

- [DOC /network/fc/wwpn-aliases](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	
alias	string	path	True	

### 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.

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an FC WWPN alias

```
GET /network/fc/wwpn-aliases/{svm.uuid}/{alias}
```

Retrieves an FC WWPN alias.

### Related ONTAP commands

- `vserver fcp wwpn-alias show`

### Learn more

- [DOC /network/fc/wwpn-aliases](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	The unique identifier of the SVM in which the alias is found.
alias	string	path	True	The name of FC WWPN alias.

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

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
alias	string	The FC WWPN alias. Required in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
wwpn	string	The FC initiator WWPN. Required in POST.

### Example response

```
{
  "_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.

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

# Manage FC services for SVMs

## 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.

### 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 the 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'
```

## Retrieve FC protocol services

GET /protocols/san/fcp/services

Retrieves FC Protocol services.

### Related ONTAP commands

- `vserver fcp show`

### Learn more

- [DOC /protocols/san/fcp/services](#)

### Parameters

Name	Type	In	Required	Description
enabled	boolean	query	False	Filter by enabled
target.name	string	query	False	Filter by target.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.

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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">fcp_service</a> ]	



## 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": {
      "name": "20:00:00:50:56:bb:b2:4b"
    }
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

target

Name	Type	Description
name	string	<p>The target name of the FC Protocol service. This is generated for the SVM during POST.</p> <p>The FC Protocol target name is a world wide node name (WWNN).</p> <p>If required, the target name can be modified using the ONTAP command line.</p> <ul style="list-style-type: none"> <li>• example: 20:00:00:50:56:bb:b2:4b</li> <li>• maxLength: 128</li> <li>• minLength: 1</li> <li>• readOnly: 1</li> </ul>

#### 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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	<p>The administrative state of the FC Protocol service. The FC Protocol service can be disabled to block all FC Protocol connectivity to the SVM.</p> <p>This is optional in POST and PATCH. The default setting is <i>true</i> (enabled) in POST.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">target</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Create an FC protocol service

POST `/protocols/san/fcp/services`

Creates an FC Protocol service.

#### Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the FC service.

#### Related ONTAP commands

- `vserver fcp create`

#### Learn more

- [DOC /protocols/san/fcp/services](#)

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
enabled	boolean	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 <i>true</i> (enabled) in POST.
svm	svm	SVM, applies only to SVM-scoped objects.
target	target	

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "name": "20:00:00:50:56:bb:b2:4b"
  }
}
```

### Response

```
Status: 201, Created
```

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">fcp_service</a> ]	

### 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": {
      "name": "20:00:00:50:56:bb:b2:4b"
    }
  }
}

```

### Error

Status: Default

Error Code	Description
1115127	The cluster lacks a valid FCP license.
2621462	The supplied 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.
2621507	The Fibre Channel Protocol is not allowed for the specified SVM.
5374092	The Fibre Channel Protocol 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.
5374082	The Fibre Channel Protocol service already exists 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

target

Name	Type	Description
name	string	<p>The target name of the FC Protocol service. This is generated for the SVM during POST.</p> <p>The FC Protocol target name is a world wide node name (WWNN).</p> <p>If required, the target name can be modified using the ONTAP command line.</p> <ul style="list-style-type: none"><li>• example: 20:00:00:50:56:bb:b2:4b</li><li>• maxLength: 128</li><li>• minLength: 1</li><li>• readOnly: 1</li></ul>

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	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 <i>true</i> (enabled) in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">target</a>	

[\\_links](#)

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

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

### Delete an FC protocol service

```
DELETE /protocols/san/fcp/services/{svm.uuid}
```

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

- [DOC /protocols/san/fcp/services](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	

### 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.
5373960	The Fibre Channel Protocol service cannot be removed while it is enabled.

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an FC protocol service

GET /protocols/san/fcp/services/{svm.uuid}

Retrieves a Fibre Channel Protocol service.

## Related ONTAP commands

- `vserver fcp show`

## Learn more

- [DOC /protocols/san/fcp/services](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	The unique identifier of the SVM for which to retrieve the FC Protocol service.
fields	array[string]	query	False	Specify the fields to return.

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	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 <i>true</i> (enabled) in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">target</a>	

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "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.

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

target

Name	Type	Description
name	string	<p>The target name of the FC Protocol service. This is generated for the SVM during POST.</p> <p>The FC Protocol target name is a world wide node name (WWNN).</p> <p>If required, the target name can be modified using the ONTAP command line.</p> <ul style="list-style-type: none"><li>• example: 20:00:00:50:56:bb:b2:4b</li><li>• maxLength: 128</li><li>• minLength: 1</li><li>• readOnly: 1</li></ul>

error\_arguments



Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Update an FC protocol service

PATCH /protocols/san/fcp/services/{svm.uuid}

Updates an FC Protocol service.

#### Related ONTAP commands

- `vserver fcp modify`
- `vserver fcp start`
- `vserver fcp stop`

#### Learn more

- [DOC /protocols/san/fcp/services](#)

#### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	The unique identifier of the SVM whose FC Protocol service is to be patched.

#### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	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 <i>true</i> (enabled) in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">target</a>	

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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.

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

target

Name	Type	Description
name	string	<p>The target name of the FC Protocol service. This is generated for the SVM during POST.</p> <p>The FC Protocol target name is a world wide node name (WWNN).</p> <p>If required, the target name can be modified using the ONTAP command line.</p> <ul style="list-style-type: none"><li>• example: 20:00:00:50:56:bb:b2:4b</li><li>• maxLength: 128</li><li>• minLength: 1</li><li>• readOnly: 1</li></ul>

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	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 <i>true</i> (enabled) in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">target</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage SAN igroups

## 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 (*fc*), 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 [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": "svml",
        "_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 the 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 disallowed. 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 disallowed. 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'
```

## Retrieve initiator groups

GET /protocols/san/igroups

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 [DOC Requesting specific fields](#) to learn more.

- `lun_maps.*`

### Related ONTAP commands

- `lun igroup show`
- `lun mapping show`

### Learn more

- [DOC /protocols/san/igroups](#)

### Parameters

<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
protocol	string	query	False	Filter by protocol
delete_on_unmap	boolean	query	False	Filter by delete_on_unmap
initiators.igroup.uuid	string	query	False	Filter by initiators.igroup.uuid
initiators.name	string	query	False	Filter by initiators.name
os_type	string	query	False	Filter by os_type
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
lun_maps.logical_unit_number	integer	query	False	Filter by lun_maps.logical_unit_number
lun_maps.lun.uuid	string	query	False	Filter by lun_maps.lun.uuid
lun_maps.lun.name	string	query	False	Filter by lun_maps.lun.name
lun_maps.lun.node.name	string	query	False	Filter by lun_maps.lun.node.name
lun_maps.lun.node.uuid	string	query	False	Filter by lun_maps.lun.node.uuid
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.

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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">igroup</a> ]	

## Example response

```
{
  "_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": "aix",
  "protocol": "fc",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "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
next	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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.

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	The unique identifier of the initiator group.

igroup\_initiator\_no\_records

Name	Type	Description
_links	<a href="#">_links</a>	
igroup	<a href="#">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	Type	Description
name	string	<p>The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the <code>records</code> property is used.</p> <p>An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is <i>iqn.yyyy-mm.reverse_domain_name:any</i>. The iSCSI EUI format consists of the <i>eui.</i> prefix followed by 16 hexadecimal characters.</p>

node

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

lun

The LUN to which the initiator group is mapped.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the LUN.
node	<a href="#">node</a>	
uuid	string	The unique identifier of the LUN.

lun\_maps

A LUN map with which the initiator group is associated.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
logical_unit_number	integer	The logical unit number assigned to the LUN for initiators in the initiator group.

Name	Type	Description
lun	<a href="#">lun</a>	The LUN to which the initiator group is mapped.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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 (*fc*), 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 [DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}](#) for more details.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
delete_on_unmap	boolean	An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in PATCH only; not available in POST. This property defaults to <i>false</i> when the initiator group is created.

Name	Type	Description
initiators	array[ <a href="#">igroup_initiator_no_records</a> ]	<p>The initiators that are members of the group. Optional in POST.</p> <p>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 <code>/protocols/san/igroups/{igroup.uuid}/initiators</code> endpoint. See <a href="#">DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}</a> for more details.</p>
lun_maps	array[ <a href="#">lun_maps</a> ]	<p>All LUN maps with which the initiator is associated.</p> <p>There is an added cost to retrieving property values for <code>lun_maps</code>. They not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
name	string	<p>The name of the initiator group. Required in POST; optional in PATCH.</p> <p>Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.</p>
os_type	string	<p>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.</p>

Name	Type	Description
protocol	string	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 <i>mixed</i> .  The protocol of an initiator group cannot be changed after creation of the group.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the initiator group.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create an initiator group

POST /protocols/san/igroups

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

- [DOC /protocols/san/igroups](#)

## Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>delete_on_unmap</code>	boolean	An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in PATCH only; not available in POST. This property defaults to <i>false</i> when the initiator group is created.
<code>initiators</code>	array[ <a href="#">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 <code>/protocols/san/igroups/{igroup.uuid}/initiators</code> endpoint. See <a href="#">DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}</a> for more details.

Name	Type	Description
lun_maps	array[lun_maps]	<p>All LUN maps with which the initiator is associated.</p> <p>There is an added cost to retrieving property values for <code>lun_maps</code>. They not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
name	string	<p>The name of the initiator group. Required in POST; optional in PATCH.</p> <p>Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.</p>
os_type	string	<p>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.</p>
protocol	string	<p>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 <i>mixed</i>.</p> <p>The protocol of an initiator group cannot be changed after creation of the group.</p>
svm	svm	<p>SVM, applies only to SVM-scoped objects.</p>
uuid	string	<p>The unique identifier of the initiator group.</p>

## Example request

```
{
  "_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": "aix",
"protocol": "fcp",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"name": "svm1",
"uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}

```

#### Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">igroup</a> ]	



## Example response

```
{
  "_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-aelc-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": "aix",
"protocol": "fcp",
"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
2621462	The supplied 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.
5374038	An invalid Fibre Channel WWPN was supplied.

Error Code	Description
5374039	An invalid iSCSI initiator name was supplied.
5373992	A supplied initiator name was too long to be valid.
5373993	A supplied initiator name did not match any valid format.
5373978	A supplied initiator name looks like an iSCSI EUI initiator, but the format is invalid.
5373977	A supplied initiator name looks like an iSCSI EUI initiator, but the length is invalid.
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.
5373969	A supplied initiator name looks like an iSCSI IQN initiator, but the portions after the prefix are missing.
5374023	An initiator group with the same name already exists.
5373958	An invalid initiator group name was supplied.
5374732	An initiator is already in another initiator group with a conflicting operating system type.
5373966	An initiator group cannot be created in an SVM is configured for NVMe.

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	<a href="#">href</a>	

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.

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	The unique identifier of the initiator group.

igroup\_initiator\_no\_records

Name	Type	Description
_links	<a href="#">_links</a>	
igroup	<a href="#">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	Type	Description
name	string	<p>The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the <code>records</code> property is used.</p> <p>An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is <i>iqn.yyyy-mm.reverse_domain_name:any</i>. The iSCSI EUI format consists of the <i>eui.</i> prefix followed by 16 hexadecimal characters.</p>

node

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

lun

The LUN to which the initiator group is mapped.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the LUN.
node	<a href="#">node</a>	
uuid	string	The unique identifier of the LUN.

lun\_maps

A LUN map with which the initiator group is associated.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
logical_unit_number	integer	The logical unit number assigned to the LUN for initiators in the initiator group.

Name	Type	Description
lun	<a href="#">lun</a>	The LUN to which the initiator group is mapped.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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 (*fc*), 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 [DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}](#) for more details.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
delete_on_unmap	boolean	An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in PATCH only; not available in POST. This property defaults to <i>false</i> when the initiator group is created.

Name	Type	Description
initiators	array[ <a href="#">igroup_initiator_no_records</a> ]	<p>The initiators that are members of the group. Optional in POST.</p> <p>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 <code>/protocols/san/igroups/{igroup.uuid}/initiators</code> endpoint. See <a href="#">DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}</a> for more details.</p>
lun_maps	array[ <a href="#">lun_maps</a> ]	<p>All LUN maps with which the initiator is associated.</p> <p>There is an added cost to retrieving property values for <code>lun_maps</code>. They not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
name	string	<p>The name of the initiator group. Required in POST; optional in PATCH.</p> <p>Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.</p>
os_type	string	<p>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.</p>



Name	Type	Description
protocol	string	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 <i>mixed</i> .  The protocol of an initiator group cannot be changed after creation of the group.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the initiator group.

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve initiators of an initiator group

GET /protocols/san/igroups/{igroup.uuid}/initiators

Retrieves initiators of an initiator group.

### Related ONTAP commands

- `lun igroup show`

### Learn more

- [DOC /protocols/san/igroups](#)

### Parameters

Name	Type	In	Required	Description
igroup.uuid	string	path	True	The unique identifier of the initiator group.
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">igroup_initiator</a> ]	

## Example response

```
{
  "_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-ae1c-123478563412"
    },
    "name": "iqn.1998-01.com.corp.iscsi:name1",
    "records": {
      "_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
5374852	The initiator group specified in the URI does not exist.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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.

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	The unique identifier of the initiator group.

records

Name	Type	Description
_links	<a href="#">_links</a>	
igroup	<a href="#">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	Type	Description
name	string	<p>The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the <code>records</code> property is used.</p> <p>An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is <i>iqn.yyyy-mm.reverse_domain_name.any</i>. The iSCSI EUI format consists of the <i>eui.</i> prefix followed by 16 hexadecimal characters.</p>

igroup\_initiator

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
igroup	<a href="#">igroup</a>	<p>The initiator group in which the initiator is found.</p> <p>Note that this does not mean that the initiator cannot also be found in other initiator groups.</p>
name	string	<p>The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the <code>records</code> property is used.</p> <p>An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is <i>iqn.yyyy-mm.reverse_domain_name.any</i>. The iSCSI EUI format consists of the <i>eui.</i> prefix followed by 16 hexadecimal characters.</p>

Name	Type	Description
records	array[ <a href="#">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 <code>name</code> property is used.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Add initiators to an initiator group

POST `/protocols/san/igroups/{igroup.uuid}/initiators`

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

- [DOC /protocols/san/igroups](#)



## Parameters

Name	Type	In	Required	Description
igroup.uuid	string	path	True	The unique identifier of the initiator group.

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
igroup	<a href="#">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	The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the <code>records</code> 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 <i>iqn.yyyy-mm.reverse_domain_name:any</i> . The iSCSI EUI format consists of the <i>eui.</i> prefix followed by 16 hexadecimal characters.
records	array[ <a href="#">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 <code>name</code> property is used.

## Example request

```
{
  "_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",
  "records": {
    "_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"
  }
}
```

## Response

Status: 201, Created

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records.

Name	Type	Description
records	array[igroup_initiator]	

## Example response

```
{
  "_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-ae1c-123478563412"
    },
    "name": "iqn.1998-01.com.corp.iscsi:name1",
    "records": {
      "_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
5374033	Initiators must be supplied.
5374038	An invalid Fibre Channel WWPN was supplied.
5374039	An invalid iSCSI initiator name was supplied.
5374852	The initiator group specified in the URI does not exist.
5374853	You can add initiators to an initiator group using the <code>records</code> property, or the <code>name&gt;</code> property, but you cannot use both in the same request.
5374854	Only <code>records</code> property elements must be populated with only the <code>name</code> property values.
5374035	A supplied initiator is already in the initiator group.
5373992	A supplied initiator name was too long to be valid.
5373993	A supplied initiator name did not match any valid format.
5373978	A supplied initiator name looks like an iSCSI EUI initiator, but the format is invalid.
5373977	A supplied initiator name looks like an iSCSI EUI initiator, but the length is invalid.
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.
5373969	A supplied initiator name looks like an iSCSI IQN initiator, but the portions after the prefix are missing.
5374734	An initiator is already in another initiator group with a conflicting operating system type.
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.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

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.

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	The unique identifier of the initiator group.

records

Name	Type	Description
_links	<a href="#">_links</a>	
igroup	<a href="#">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	Type	Description
name	string	<p>The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the <code>records</code> property is used.</p> <p>An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is <i>iqn.yyyy-mm.reverse_domain_name.any</i>. The iSCSI EUI format consists of the <i>eui.</i> prefix followed by 16 hexadecimal characters.</p>

igroup\_initiator

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
igroup	<a href="#">igroup</a>	<p>The initiator group in which the initiator is found.</p> <p>Note that this does not mean that the initiator cannot also be found in other initiator groups.</p>
name	string	<p>The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the <code>records</code> property is used.</p> <p>An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is <i>iqn.yyyy-mm.reverse_domain_name.any</i>. The iSCSI EUI format consists of the <i>eui.</i> prefix followed by 16 hexadecimal characters.</p>



Name	Type	Description
records	array[ <a href="#">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 <code>name</code> property is used.

#### `_links`

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### `error_arguments`

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### `error`

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete an initiator from an initiator group

```
DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}
```

Deletes an initiator from an initiator group.

#### Related ONTAP commands

- `lun igroup remove`

## Learn more

- [DOC /protocols/san/igroups](#)

## Parameters

Name	Type	In	Required	Description
igroup.uuid	string	path	True	The unique identifier of the initiator group.
name	string	path	True	The initiator name.
allow_delete_while_mapped	boolean	query	False	<p>Allow deletion of an initiator from of a mapped initiator group.</p> <p>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.</p> <p><b>This parameter should be used with caution.</b></p> <ul style="list-style-type: none"><li>• Default value:</li></ul>

## Response

```
Status: 200, Ok
```

## Error

```
Status: Default
```

## ONTAP Error Response Codes

Error Code	Description
5374034	The initiator is not a member of the group.

Error Code	Description
5374852	The initiator group specified in the URI does not exist.
1254213	The initiator group is mapped to one or more LUNs and allow_delete_while_mapped has not been specified.

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an initiator

GET /protocols/san/igroups/{igroup.uuid}/initiators/{name}

Retrieves an initiator of an initiator group.

### Related ONTAP commands

- `lun igroup show`

### Learn more

- [DOC /protocols/san/igroups](#)

### Parameters

Name	Type	In	Required	Description
igroup.uuid	string	path	True	The unique identifier of the initiator group.
name	string	path	True	The initiator name.
fields	array[string]	query	False	Specify the fields to return.

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
igroup	<a href="#">igroup</a>	<p>The initiator group in which the initiator is found.</p> <p>Note that this does not mean that the initiator cannot also be found in other initiator groups.</p>
name	string	<p>The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the <code>records</code> property is used.</p> <p>An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is <i>iqn.yyyy-mm.reverse_domain_name:any</i>. The iSCSI EUI format consists of the <i>eui.</i> prefix followed by 16 hexadecimal characters.</p>
records	array[ <a href="#">records</a> ]	<p>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 <code>name</code> property is used.</p>

## Example response

```
{
  "_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",
  "records": {
    "_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.

Error Code	Description
5374852	The initiator group specified in the URI 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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.

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	The unique identifier of the initiator group.

records

Name	Type	Description
_links	<a href="#">_links</a>	
igroup	<a href="#">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	Type	Description
name	string	<p>The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the <code>records</code> property is used.</p> <p>An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is <i>iqn.yyyy-mm.reverse_domain_name.any</i>. The iSCSI EUI format consists of the <i>eui.</i> prefix followed by 16 hexadecimal characters.</p>

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete an initiator group

```
DELETE /protocols/san/igroups/{uuid}
```

Deletes an initiator group.

#### Related ONTAP commands

- `lun igroup delete`

## Learn more

- [DOC /protocols/san/igroups](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier of the initiator group.
allow_delete_while_mapped	boolean	query	False	<p>Allow deletion of a mapped initiator group.</p> <p>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.</p> <p><b>This parameter should be used with caution.</b></p> <ul style="list-style-type: none"><li>• Default value:</li></ul>

## Response

Status: 200, Ok

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
5374852	The initiator group does not exist.
1254213	The initiator group is mapped to one or more LUNs and <code>allow_delete_while_mapped</code> has not been specified.

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an initiator group

GET /protocols/san/igroups/{uuid}

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 [DOC Requesting specific fields](#) to learn more.

- `lun_maps.*`

### Related ONTAP commands

- `lun igroup show`
- `lun mapping show`

### Learn more

- [DOC /protocols/san/igroups](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier of the initiator group.
fields	array[string]	query	False	Specify the fields to return.

### Response

Status: 200, Ok

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>delete_on_unmap</code>	boolean	An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in PATCH only; not available in POST. This property defaults to <i>false</i> when the initiator group is created.

Name	Type	Description
initiators	array[ <a href="#">igroup_initiator_no_records</a> ]	<p>The initiators that are members of the group. Optional in POST.</p> <p>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 <code>/protocols/san/igroups/{igroup.uuid}/initiators</code> endpoint. See <a href="#">DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}</a> for more details.</p>
lun_maps	array[ <a href="#">lun_maps</a> ]	<p>All LUN maps with which the initiator is associated.</p> <p>There is an added cost to retrieving property values for <code>lun_maps</code>. They not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
name	string	<p>The name of the initiator group. Required in POST; optional in PATCH.</p> <p>Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.</p>
os_type	string	<p>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.</p>

Name	Type	Description
protocol	string	<p>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 <i>mixed</i>.</p> <p>The protocol of an initiator group cannot be changed after creation of the group.</p>
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the initiator group.

## Example response

```
{
  "_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": "aix",
"protocol": "fcp",
"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.

Name	Type	Description
error	<a href="#">error</a>	



## 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	<a href="#">href</a>	

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.

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	The unique identifier of the initiator group.

igroup\_initiator\_no\_records

Name	Type	Description
_links	<a href="#">_links</a>	
igroup	<a href="#">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	Type	Description
name	string	<p>The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the <code>records</code> property is used.</p> <p>An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is <i>iqn.yyyy-mm.reverse_domain_name:any</i>. The iSCSI EUI format consists of the <i>eui.</i> prefix followed by 16 hexadecimal characters.</p>

node

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

lun

The LUN to which the initiator group is mapped.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the LUN.
node	<a href="#">node</a>	
uuid	string	The unique identifier of the LUN.

lun\_maps

A LUN map with which the initiator group is associated.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
logical_unit_number	integer	The logical unit number assigned to the LUN for initiators in the initiator group.

Name	Type	Description
lun	<a href="#">lun</a>	The LUN to which the initiator group is mapped.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update an initiator group

PATCH /protocols/san/igroups/{uuid}

Updates an initiator group.

### Related ONTAP commands

- `lun igroup modify`
- `lun igroup rename`

## Learn more

- [DOC /protocols/san/igroups](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier of the initiator group.

## Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
delete_on_unmap	boolean	An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in PATCH only; not available in POST. This property defaults to <i>false</i> when the initiator group is created.
initiators	array[ <a href="#">igroup_initiator_no_records</a> ]	<p>The initiators that are members of the group. Optional in POST.</p> <p>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 <code>/protocols/san/igroups/{igroup.uuid}/initiators</code> endpoint. See <a href="#">DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}</a> for more details.</p>
lun_maps	array[ <a href="#">lun_maps</a> ]	<p>All LUN maps with which the initiator is associated.</p> <p>There is an added cost to retrieving property values for <code>lun_maps</code>. They not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
name	string	<p>The name of the initiator group. Required in POST; optional in PATCH.</p> <p>Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.</p>
os_type	string	<p>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.</p>
protocol	string	<p>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 <i>mixed</i>.</p> <p>The protocol of an initiator group cannot be changed after creation of the group.</p>
svm	svm	<p>SVM, applies only to SVM-scoped objects.</p>
uuid	string	<p>The unique identifier of the initiator group.</p>

## Example request

```
{
  "_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": "aix",
"protocol": "fc",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"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
5374852	The initiator group does not exist.
5374868	The initiator group was partially modified before an error was encountered renaming the initiator group.
5374023	A rename operation failed because an initiator group with the same name already exists.
5373958	An invalid initiator group name was supplied for a rename operation.
5374733	An initiator is already in another initiator group with a conflicting operating system type.



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	<a href="#">href</a>	

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.

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	The unique identifier of the initiator group.

igroup\_initiator\_no\_records

Name	Type	Description
_links	<a href="#">_links</a>	
igroup	<a href="#">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	Type	Description
name	string	<p>The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the <code>records</code> property is used.</p> <p>An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is <i>iqn.yyyy-mm.reverse_domain_name:any</i>. The iSCSI EUI format consists of the <i>eui.</i> prefix followed by 16 hexadecimal characters.</p>

node

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

lun

The LUN to which the initiator group is mapped.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the LUN.
node	<a href="#">node</a>	
uuid	string	The unique identifier of the LUN.

lun\_maps

A LUN map with which the initiator group is associated.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
logical_unit_number	integer	The logical unit number assigned to the LUN for initiators in the initiator group.

Name	Type	Description
lun	<a href="#">lun</a>	The LUN to which the initiator group is mapped.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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 (*fc*), 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 [DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}](#) for more details.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
delete_on_unmap	boolean	An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in PATCH only; not available in POST. This property defaults to <i>false</i> when the initiator group is created.

Name	Type	Description
initiators	array[ <a href="#">igroup_initiator_no_records</a> ]	<p>The initiators that are members of the group. Optional in POST.</p> <p>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 <code>/protocols/san/igroups/{igroup.uuid}/initiators</code> endpoint. See <a href="#">DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}</a> for more details.</p>
lun_maps	array[ <a href="#">lun_maps</a> ]	<p>All LUN maps with which the initiator is associated.</p> <p>There is an added cost to retrieving property values for <code>lun_maps</code>. They not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
name	string	<p>The name of the initiator group. Required in POST; optional in PATCH.</p> <p>Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.</p>
os_type	string	<p>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.</p>

Name	Type	Description
protocol	string	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 <i>mixed</i> .  The protocol of an initiator group cannot be changed after creation of the group.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the initiator group.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage iSCSI credentials

### 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'
```

### Retrieve iSCSI credentials

GET /protocols/san/iscsi/credentials

Retrieves iSCSI credentials.

## Related ONTAP commands

- `vserver iscsi security show`

## Learn more

- [DOC /protocols/san/iscsi/credentials](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
initiator	string	query	False	Filter by initiator
initiator_address.ranges.end	string	query	False	Filter by initiator_address.ranges.end
initiator_address.ranges.start	string	query	False	Filter by initiator_address.ranges.start
initiator_address.ranges.family	string	query	False	Filter by initiator_address.ranges.family
initiator_address.masks.address	string	query	False	Filter by initiator_address.masks.address
initiator_address.masks.netmask	string	query	False	Filter by initiator_address.masks.netmask
initiator_address.masks.family	string	query	False	Filter by initiator_address.masks.family
chap.inbound.user	string	query	False	Filter by chap.inbound.user
chap.outbound.user	string	query	False	Filter by chap.outbound.user

Name	Type	In	Required	Description
authentication_type	string	query	False	Filter by authentication_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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">iscsi_credentials</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "authentication_type": "chap",
    "initiator": "iqn.1998-01.com.corp.iscsi:namel",
    "initiator_address": {
      "masks": {
        "address": "10.10.10.7",
        "family": "ipv4",
        "netmask": "24"
      },
      "ranges": {
        "end": "10.10.10.7",
        "family": "ipv4",
        "start": "10.10.10.7"
      }
    },
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

inbound

Inbound CHAP credentials.

Name	Type	Description
password	string	The inbound CHAP password. Write-only; optional in POST and PATCH.
user	string	The inbound CHAP user name. Optional in POST and PATCH.

outbound

Output CHAP credentials.

Name	Type	Description
password	string	The outbound CHAP password. Write-only; optional in POST and PATCH.
user	string	The outbound CHAP user name. Optional in POST and PATCH.

chap

Challenge-Handshake Authentication Protocol (CHAP) credentials.

Name	Type	Description
inbound	<a href="#">inbound</a>	Inbound CHAP credentials.
outbound	<a href="#">outbound</a>	Output CHAP credentials.

ip\_info

IP information

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.

ip\_address\_range

IP address range

Name	Type	Description
end	string	IPv4 or IPv6 address
family	string	IPv4 or IPv6
start	string	IPv4 or IPv6 address

initiator\_address

Initiator address ranges.

Name	Type	Description
masks	array[ <a href="#">ip_info</a> ]	
ranges	array[ <a href="#">ip_address_range</a> ]	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### iscsi\_credentials

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
authentication_type	string	The iSCSI authentication type. Required in POST and optional in PATCH.
chap	<a href="#">chap</a>	Challenge-Handshake Authentication Protocol (CHAP) credentials.
initiator	string	The iSCSI initiator to which the credentials apply. Required in POST.
initiator_address	<a href="#">initiator_address</a>	Initiator address ranges.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

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

## Create iSCSI credentials

POST /protocols/san/iscsi/credentials

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

- [DOC /protocols/san/iscsi/credentials](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>authentication_type</code>	string	The iSCSI authentication type. Required in POST and optional in PATCH.
<code>chap</code>	<a href="#">chap</a>	Challenge-Handshake Authentication Protocol (CHAP) credentials.

Name	Type	Description
initiator	string	The iSCSI initiator to which the credentials apply. Required in POST.
initiator_address	<a href="#">initiator_address</a>	Initiator address ranges.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_type": "chap",
  "initiator": "iqn.1998-01.com.corp.iscsi:name1",
  "initiator_address": {
    "masks": {
      "address": "10.10.10.7",
      "family": "ipv4",
      "netmask": "24"
    },
    "ranges": {
      "end": "10.10.10.7",
      "family": "ipv4",
      "start": "10.10.10.7"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">iscsi_credentials</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "authentication_type": "chap",
    "initiator": "iqn.1998-01.com.corp.iscsi:namel",
    "initiator_address": {
      "masks": {
        "address": "10.10.10.7",
        "family": "ipv4",
        "netmask": "24"
      },
      "ranges": {
        "end": "10.10.10.7",
        "family": "ipv4",
        "start": "10.10.10.7"
      }
    },
    "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	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 <code>svm.name</code> or <code>svm.uuid</code> must be supplied.
5373969	A non-empty qualifier is required after the prefix. An example of a valid IQN is <code>iqn.1995-08.com.example:string</code> .
5373970	The IQN prefix is invalid. The correct IQN prefix is <code>iqn</code> . An example of a valid IQN is <code>iqn.1995-08.com.example:string</code> .
5373971	The date field is invalid. A valid date field is <code>yyyy-mm</code> . An example of a valid IQN is <code>iqn.1995-08.com.example:string</code> .
5373972	The naming authority and string fields can contain only the characters <code>a-z</code> , <code>0-9</code> , <code>.</code> , <code>:</code> , and <code>-</code> .
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 <code>eui.XXXXXXXXXXXXXXXXXX</code> , where X is a hexadecimal digit.
5373997	The initiator name supplied is invalid. The valid initiator name formats are <code>iqn.1995-08.com.example:string</code> or <code>eui.0123456789abcdef</code> .
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.
5374900	Setting the CHAP authentication properties are not supported with authentication types <code>none</code> or <code>deny</code> .
5374147	The CHAP inbound and outbound passwords must be different.
5374149	The inbound user and password properties are required for CHAP authentication.

Error Code	Description
5374150	Outbound CHAP authentication requires an outbound password.
5374855	The value for property <code>initiator_address.ranges.start</code> is greater than the value for property <code>initiator_address.ranges.end</code> .
5374856	The value for property <code>initiator_address.ranges.start</code> does not belong to the same IP address family as the value for property <code>initiator_address.ranges.end</code> .

Name	Type	Description
error	<a href="#">error</a>	

### 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	<a href="#">href</a>	

inbound

Inbound CHAP credentials.

Name	Type	Description
password	string	The inbound CHAP password. Write-only; optional in POST and PATCH.
user	string	The inbound CHAP user name. Optional in POST and PATCH.

outbound

Output CHAP credentials.

Name	Type	Description
password	string	The outbound CHAP password. Write-only; optional in POST and PATCH.
user	string	The outbound CHAP user name. Optional in POST and PATCH.

chap

Challenge-Handshake Authentication Protocol (CHAP) credentials.

Name	Type	Description
inbound	<a href="#">inbound</a>	Inbound CHAP credentials.
outbound	<a href="#">outbound</a>	Output CHAP credentials.

ip\_info

## IP information

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.

## ip\_address\_range

### IP address range

Name	Type	Description
end	string	IPv4 or IPv6 address
family	string	IPv4 or IPv6
start	string	IPv4 or IPv6 address

## initiator\_address

### Initiator address ranges.

Name	Type	Description
masks	array[ip_info]	
ranges	array[ip_address_range]	

## svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

## iscsi\_credentials

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
authentication_type	string	The iSCSI authentication type. Required in POST and optional in PATCH.
chap	<a href="#">chap</a>	Challenge-Handshake Authentication Protocol (CHAP) credentials.
initiator	string	The iSCSI initiator to which the credentials apply. Required in POST.
initiator_address	<a href="#">initiator_address</a>	Initiator address ranges.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

[\\_links](#)

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete iSCSI credentials

```
DELETE /protocols/san/iscsi/credentials/{svm.uuid}/{initiator}
```

Deletes specified iSCSI credentials.

### Related ONTAP commands

- `vserver iscsi security delete`

### Learn more

- [DOC /protocols/san/iscsi/credentials](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	
initiator	string	path	True	

### 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 <code>svm.name</code> or <code>svm.uuid</code> 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.

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve specific iSCSI credentials

GET /protocols/san/iscsi/credentials/{svm.uuid}/{initiator}

Retrieves specified iSCSI credentials.

## Related ONTAP commands

- `vserver iscsi security show`

## Learn more

- [DOC /protocols/san/iscsi/credentials](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	The unique identifier of an SVM.
initiator	string	path	True	The iSCSI initiator of the credentials object.
fields	array[string]	query	False	Specify the fields to return.

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
authentication_type	string	The iSCSI authentication type. Required in POST and optional in PATCH.
chap	<a href="#">chap</a>	Challenge-Handshake Authentication Protocol (CHAP) credentials.
initiator	string	The iSCSI initiator to which the credentials apply. Required in POST.
initiator_address	<a href="#">initiator_address</a>	Initiator address ranges.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.



## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_type": "chap",
  "initiator": "iqn.1998-01.com.corp.iscsi:name1",
  "initiator_address": {
    "masks": {
      "address": "10.10.10.7",
      "family": "ipv4",
      "netmask": "24"
    },
    "ranges": {
      "end": "10.10.10.7",
      "family": "ipv4",
      "start": "10.10.10.7"
    }
  },
  "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	<a href="#">href</a>	

inbound

Inbound CHAP credentials.

Name	Type	Description
password	string	The inbound CHAP password. Write-only; optional in POST and PATCH.
user	string	The inbound CHAP user name. Optional in POST and PATCH.

outbound

Output CHAP credentials.

Name	Type	Description
password	string	The outbound CHAP password. Write-only; optional in POST and PATCH.
user	string	The outbound CHAP user name. Optional in POST and PATCH.

chap

Challenge-Handshake Authentication Protocol (CHAP) credentials.

Name	Type	Description
inbound	<a href="#">inbound</a>	Inbound CHAP credentials.
outbound	<a href="#">outbound</a>	Output CHAP credentials.

ip\_info

## IP information

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.

## ip\_address\_range

### IP address range

Name	Type	Description
end	string	IPv4 or IPv6 address
family	string	IPv4 or IPv6
start	string	IPv4 or IPv6 address

## initiator\_address

### Initiator address ranges.

Name	Type	Description
masks	array[ip_info]	
ranges	array[ip_address_range]	

## svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update iSCSI credentials

```
PATCH /protocols/san/iscsi/credentials/{svm.uuid}/{initiator}
```

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

- [DOC /protocols/san/iscsi/credentials](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	The unique identifier of an SVM.
initiator	string	path	True	The iSCSI initiator of the credentials object.

Name	Type	In	Required	Description
add_initiator_addresses	boolean	query	False	If <i>true</i> , 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	If <i>true</i> , 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

Name	Type	Description
_links	<a href="#">_links</a>	
authentication_type	string	The iSCSI authentication type. Required in POST and optional in PATCH.
chap	<a href="#">chap</a>	Challenge-Handshake Authentication Protocol (CHAP) credentials.
initiator	string	The iSCSI initiator to which the credentials apply. Required in POST.
initiator_address	<a href="#">initiator_address</a>	Initiator address ranges.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_type": "chap",
  "initiator": "iqn.1998-01.com.corp.iscsi:name1",
  "initiator_address": {
    "masks": {
      "address": "10.10.10.7",
      "family": "ipv4",
      "netmask": "24"
    },
    "ranges": {
      "end": "10.10.10.7",
      "family": "ipv4",
      "start": "10.10.10.7"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "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 <code>svm.name</code> or <code>svm.uuid</code> must be supplied.
5374145	The iSCSI security password must contain an even number of valid hex digits.
5374900	Setting the CHAP authentication properties are not supported with authentication types <i>none</i> or <i>deny</i> .
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 <code>initiator_address.ranges.start</code> is greater than the value for property <code>initiator_address.ranges.end</code> .
5374856	The value for property <code>initiator_address.ranges.start</code> does not belong to the same IP address family as the value for property <code>initiator_address.ranges.end</code> .
5374895	The iSCSI security credential does not exist on the specified SVM.

Name	Type	Description
error	<a href="#">error</a>	



## 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	<a href="#">href</a>	

inbound

Inbound CHAP credentials.

Name	Type	Description
password	string	The inbound CHAP password. Write-only; optional in POST and PATCH.
user	string	The inbound CHAP user name. Optional in POST and PATCH.

outbound

Output CHAP credentials.

Name	Type	Description
password	string	The outbound CHAP password. Write-only; optional in POST and PATCH.
user	string	The outbound CHAP user name. Optional in POST and PATCH.

chap

Challenge-Handshake Authentication Protocol (CHAP) credentials.

Name	Type	Description
inbound	<a href="#">inbound</a>	Inbound CHAP credentials.
outbound	<a href="#">outbound</a>	Output CHAP credentials.

ip\_info

## IP information

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.

## ip\_address\_range

### IP address range

Name	Type	Description
end	string	IPv4 or IPv6 address
family	string	IPv4 or IPv6
start	string	IPv4 or IPv6 address

## initiator\_address

### Initiator address ranges.

Name	Type	Description
masks	array[ip_info]	
ranges	array[ip_address_range]	

## svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

## iscsi\_credentials

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
authentication_type	string	The iSCSI authentication type. Required in POST and optional in PATCH.
chap	<a href="#">chap</a>	Challenge-Handshake Authentication Protocol (CHAP) credentials.
initiator	string	The iSCSI initiator to which the credentials apply. Required in POST.
initiator_address	<a href="#">initiator_address</a>	Initiator address ranges.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage iSCSI services

### 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.

## 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": "svm1",
    "_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'
```

## Retrieve iSCSI services

GET /protocols/san/iscsi/services

Retrieves iSCSI services.

### Related ONTAP commands

- `vserver iscsi show`

### Learn more

- [DOC /protocols/san/iscsi/services](#)

### Parameters

Name	Type	In	Required	Description
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
target.name	string	query	False	Filter by target.name
target.alias	string	query	False	Filter by target.alias
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.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">iscsi_service</a> ]	

## 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": {
      "alias": "svm1",
      "name": "iqn.1992-
08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2"
    }
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

target

Name	Type	Description
alias	string	<p>The iSCSI target alias of the iSCSI service.</p> <p>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.</p> <p>Optional in POST and PATCH. In POST, this defaults to the name of the SVM.</p>

Name	Type	Description
name	string	<p>The iSCSI target name of the iSCSI service. This is generated for the SVM during POST.</p> <p>If required, the target name can be modified using the ONTAP command line.</p> <ul style="list-style-type: none"> <li>• example: iqn.1992-08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2</li> <li>• maxLength: 128</li> <li>• minLength: 1</li> <li>• readOnly: 1</li> </ul>

#### 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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	<p>The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.</p> <p>Optional in POST and PATCH. The default setting is <i>true</i> (enabled) in POST.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">target</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument



error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Create an iSCSI service

POST `/protocols/san/iscsi/services`

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

- [DOC /protocols/san/iscsi/services](#)

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>enabled</code>	boolean	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 <i>true</i> (enabled) in POST.
<code>svm</code>	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
<code>target</code>	<a href="#">target</a>	

## Example request

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

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">iscsi_service</a> ]	

## 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": {
      "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.

Error Code	Description
2621507	The iSCSI protocol is not allowed for the specified SVM.
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.
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.
5373966	An iSCSI service cannot be created in an SVM that is configured for NVMe.

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

target

Name	Type	Description
alias	string	<p>The iSCSI target alias of the iSCSI service.</p> <p>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.</p> <p>Optional in POST and PATCH. In POST, this defaults to the name of the SVM.</p>

Name	Type	Description
name	string	<p>The iSCSI target name of the iSCSI service. This is generated for the SVM during POST.</p> <p>If required, the target name can be modified using the ONTAP command line.</p> <ul style="list-style-type: none"> <li>example: iqn.1992-08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2</li> <li>maxLength: 128</li> <li>minLength: 1</li> <li>readOnly: 1</li> </ul>

#### 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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	<p>The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.</p> <p>Optional in POST and PATCH. The default setting is <i>true</i> (enabled) in POST.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">target</a>	

#### [\\_links](#)

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

## error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

## error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete an iSCSI service

DELETE /protocols/san/iscsi/services/{svm.uuid}

Deletes an iSCSI service. An iSCSI service must be disabled before it can be deleted.

### Related ONTAP commands

- `vserver iscsi delete`

### Learn more

- [DOC /protocols/san/iscsi/services](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	

### 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.

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an iSCSI service

GET /protocols/san/iscsi/services/{svm.uuid}

Retrieves an iSCSI service.

### Related ONTAP commands

- `vserver iscsi show`

### Learn more

- [DOC /protocols/san/iscsi/services](#)

### Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	The unique identifier of the SVM for which to retrieve the iSCSI service.
fields	array[string]	query	False	Specify the fields to return.

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	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 <i>true</i> (enabled) in POST.
<a href="#">svm</a>	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
<a href="#">target</a>	<a href="#">target</a>	

## Example response

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

## 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.

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

target

Name	Type	Description
alias	string	<p>The iSCSI target alias of the iSCSI service.</p> <p>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.</p> <p>Optional in POST and PATCH. In POST, this defaults to the name of the SVM.</p>

Name	Type	Description
name	string	<p>The iSCSI target name of the iSCSI service. This is generated for the SVM during POST.</p> <p>If required, the target name can be modified using the ONTAP command line.</p> <ul style="list-style-type: none"> <li>• example: iqn.1992-08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2</li> <li>• maxLength: 128</li> <li>• minLength: 1</li> <li>• readOnly: 1</li> </ul>

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Update an iSCSI service

PATCH /protocols/san/iscsi/services/{svm.uuid}

Updates an iSCSI service.

## Related ONTAP commands

- `vserver iscsi modify`
- `vserver iscsi start`
- `vserver iscsi stop`

## Learn more

- [DOC /protocols/san/iscsi/services](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	The unique identifier of the SVM for which to update the iSCSI service.

## Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	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 <i>true</i> (enabled) in POST.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">target</a>	

## Example request

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

## 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.

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

target

Name	Type	Description
alias	string	<p>The iSCSI target alias of the iSCSI service.</p> <p>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.</p> <p>Optional in POST and PATCH. In POST, this defaults to the name of the SVM.</p>

Name	Type	Description
name	string	<p>The iSCSI target name of the iSCSI service. This is generated for the SVM during POST.</p> <p>If required, the target name can be modified using the ONTAP command line.</p> <ul style="list-style-type: none"> <li>example: iqn.1992-08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2</li> <li>maxLength: 128</li> <li>minLength: 1</li> <li>readOnly: 1</li> </ul>

#### 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.

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	<p>The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.</p> <p>Optional in POST and PATCH. The default setting is <i>true</i> (enabled) in POST.</p>
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target	<a href="#">target</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## View iSCSI sessions

### 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"
    }
  }
}

```

### Retrieve iSCSI sessions

GET /protocols/san/iscsi/sessions

Retrieves iSCSI sessions.

### Related ONTAP commands

- `vserver iscsi connection show`
- `vserver iscsi session parameter show`
- `vserver iscsi session show`

## Learn more

- [DOC /protocols/san/iscsi/sessions](#)

## Parameters

Name	Type	In	Required	Description
initiator.name	string	query	False	Filter by initiator.name
initiator.alias	string	query	False	Filter by initiator.alias
target_portal_group_tag	integer	query	False	Filter by target_portal_group_tag
igroups.name	string	query	False	Filter by igroups.name
igroups.uuid	string	query	False	Filter by igroups.uuid
isid	string	query	False	Filter by isid
tsih	integer	query	False	Filter by tsih
target_portal_group	string	query	False	Filter by target_portal_group
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
connections.cid	integer	query	False	Filter by connections.cid
connections.interface.name	string	query	False	Filter by connections.interface.name
connections.interface.uuid	string	query	False	Filter by connections.interface.uuid
connections.interface.ip.address	string	query	False	Filter by connections.interface.ip.address



Name	Type	In	Required	Description
connections.interface.ip.port	integer	query	False	Filter by connections.interface.ip.port
connections.initiator_address.port	integer	query	False	Filter by connections.initiator_address.port
connections.initiator_address.address	string	query	False	Filter by connections.initiator_address.address
connections.authentication_type	string	query	False	Filter by connections.authentication_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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">iscsi_session</a> ]	

## Example response

```
{
  "_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": "chap",
      "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

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

initiator\_address

The TCP socket information for the initiator end of the connection. This is useful for network packet debugging.

Name	Type	Description
address	string	The TCP IPv4 or IPv6 address of the initiator end of the iSCSI connection.
port	integer	The TCP port number of the initiator end of the iSCSI connection.

ip

The IP information. ONTAP only supports port 3260.

Name	Type	Description
address	string	IPv4 or IPv6 address
port	integer	The TCP port number of the iSCSI access endpoint.

interface

The network interface information for the target end of the connection.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
ip	ip	The IP information. ONTAP only supports port 3260. <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

#### iscsi\_connection

An active iSCSI connection.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
authentication_type	string	The iSCSI authentication type used to establish the connection.
cid	integer	The identifier of the connection within the session.
initiator_address	<a href="#">initiator_address</a>	The TCP socket information for the initiator end of the connection. This is useful for network packet debugging.
interface	<a href="#">interface</a>	The network interface information for the target end of the connection.

#### igroups

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the initiator group.
uuid	string	The unique identifier of the initiator group.

#### initiator

The initiator that created the session.

Name	Type	Description
alias	string	The initiator alias.
name	string	The world wide unique name of the initiator.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
connections	array[ <a href="#">iscsi_connection</a> ]	The iSCSI connections that make up the iSCSI session.
igroups	array[ <a href="#">igroups</a> ]	The initiator groups in which the initiator is a member.
initiator	<a href="#">initiator</a>	The initiator that created the session.
isid	string	The initiator portion of the session identifier specified by the initiator during login.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target_portal_group	string	The target portal group to which the session belongs.



Name	Type	Description
target_portal_group_tag	integer	The target portal group tag of the session.
tsih	integer	The target session identifier handle (TSIH) of the session.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an iSCSI session

GET /protocols/san/iscsi/sessions/{svm.uuid}/{tpgroup}/{tsih}

Retrieves an iSCSI session.

### Related ONTAP commands

- `vserver iscsi connection show`
- `vserver iscsi session parameter show`
- `vserver iscsi session show`

### Learn more

- [DOC /protocols/san/iscsi/sessions](#)

## Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	The unique identifier of the SVM of the iSCSI session.
tpgroup	string	path	True	The target portal group of the iSCSI session.
tsih	integer	path	True	The target session identifying handle.
fields	array[string]	query	False	Specify the fields to return.

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
connections	array[ <a href="#">iscsi_connection</a> ]	The iSCSI connections that make up the iSCSI session.
igroups	array[ <a href="#">igroups</a> ]	The initiator groups in which the initiator is a member.
initiator	<a href="#">initiator</a>	The initiator that created the session.
isid	string	The initiator portion of the session identifier specified by the initiator during login.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
target_portal_group	string	The target portal group to which the session belongs.
target_portal_group_tag	integer	The target portal group tag of the session.

<b>Name</b>	<b>Type</b>	<b>Description</b>
tsih	integer	The target session identifier handle (TSIH) of the session.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "connections": {
    "_links": {
      "next": {
        "href": "/api/resourcelink"
      },
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "authentication_type": "chap",
    "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.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

initiator\_address

The TCP socket information for the initiator end of the connection. This is useful for network packet debugging.

Name	Type	Description
address	string	The TCP IPv4 or IPv6 address of the initiator end of the iSCSI connection.
port	integer	The TCP port number of the initiator end of the iSCSI connection.

ip

The IP information. ONTAP only supports port 3260.

Name	Type	Description
address	string	IPv4 or IPv6 address
port	integer	The TCP port number of the iSCSI access endpoint.

interface

The network interface information for the target end of the connection.

Name	Type	Description
_links	<a href="#">_links</a>	
ip	ip	The IP information. ONTAP only supports port 3260. <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

#### iscsi\_connection

An active iSCSI connection.

Name	Type	Description
_links	<a href="#">_links</a>	
authentication_type	string	The iSCSI authentication type used to establish the connection.
cid	integer	The identifier of the connection within the session.
initiator_address	<a href="#">initiator_address</a>	The TCP socket information for the initiator end of the connection. This is useful for network packet debugging.
interface	<a href="#">interface</a>	The network interface information for the target end of the connection.

#### igroups

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the initiator group.
uuid	string	The unique identifier of the initiator group.

#### initiator

The initiator that created the session.



Name	Type	Description
alias	string	The initiator alias.
name	string	The world wide unique name of the initiator.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage SAN LUN maps

### 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": "svml" }, "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": "svml",
        "_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'
```

## Retrieve LUN maps

GET /protocols/san/lun-maps

Retrieves LUN maps.

### Related ONTAP commands

- lun mapping show
- [DOC /protocols/san/lun-maps](#)

### Learn more

- [DOC /protocols/san/lun-maps](#)

### Parameters

Name	Type	In	Required	Description
logical_unit_number	integer	query	False	Filter by logical_unit_number
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
lun.uuid	string	query	False	Filter by lun.uuid
lun.node.uuid	string	query	False	Filter by lun.node.uuid
lun.node.name	string	query	False	Filter by lun.node.name
lun.name	string	query	False	Filter by lun.name
igroup.initiators	string	query	False	Filter by igroup.initiators

Name	Type	In	Required	Description
igroup.os_type	string	query	False	Filter by igroup.os_type
igroup.name	string	query	False	Filter by igroup.name
igroup.protocol	string	query	False	Filter by igroup.protocol
igroup.uuid	string	query	False	Filter by igroup.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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">lun_map</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "igroup": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "initiators": {
      },
      "name": "igroup1",
      "os_type": "aix",
      "protocol": "fc",
      "uuid": "1ad8544d-8cd1-91e0-9e1c-723478563412"
    },
    "logical_unit_number": 1,
    "lun": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "/vol/volume1/qtreen1/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

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

igroup

The initiator group to which the LUN is mapped. Required in POST by supplying either the `igroup.uuid`, `igroup.name`, or both.

Name	Type	Description
_links	<a href="#">_links</a>	
initiators	array[string]	The initiators that are members of the initiator group.
name	string	The name of the initiator group. Valid in POST.
os_type	string	The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system.
protocol	string	The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group.
uuid	string	The unique identifier of the initiator group. Valid in POST.

node

The LUN node.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name the LUN's node.
uuid	string	The unique identifier of the LUN node.

#### lun

The LUN to which the initiator group is mapped. Required in POST by supplying either the `lun.uuid`, `lun.name`, or both.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	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	<a href="#">node</a>	The LUN node.
uuid	string	The unique identifier of the LUN. Valid in POST.

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### 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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
igroup	<a href="#">igroup</a>	The initiator group to which the LUN is mapped. Required in POST by supplying either the <code>igroup.uuid</code> , <code>igroup.name</code> , or both.
logical_unit_number	integer	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. <ul style="list-style-type: none"> <li>• example: 1</li> <li>• max<b>Value</b>: 4095</li> <li>• min<b>Value</b>: 0</li> <li>• readCreate: 1</li> </ul>
lun	<a href="#">lun</a>	The LUN to which the initiator group is mapped. Required in POST by supplying either the <code>lun.uuid</code> , <code>lun.name</code> , or both.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code

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

## Create a LUN map

POST /protocols/san/lun-maps

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

- [DOC /protocols/san/lun-maps](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>igroup</code>	<a href="#">igroup</a>	The initiator group to which the LUN is mapped. Required in POST by supplying either the <code>igroup.uuid</code> , <code>igroup.name</code> , or both.

Name	Type	Description
logical_unit_number	integer	<p>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.</p> <ul style="list-style-type: none"> <li>• example: 1</li> <li>• maxValue: 4095</li> <li>• minValue: 0</li> <li>• readCreate: 1</li> </ul>
lun	lun	<p>The LUN to which the initiator group is mapped. Required in POST by supplying either the <code>lun.uuid</code>, <code>lun.name</code>, or both.</p>
svm	svm	<p>SVM, applies only to SVM-scoped objects.</p>

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "igroup": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "initiators": {
    },
    "name": "igroup1",
    "os_type": "aix",
    "protocol": "fc",
    "uuid": "1ad8544d-8cd1-91e0-9e1c-723478563412"
  },
  "logical_unit_number": 1,
  "lun": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "/vol/volume1/mtree1/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"  
}
```

### Response

Status: 201, Created

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">lun_map</a> ]	



## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "igroup": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "initiators": {
      },
      "name": "igroup1",
      "os_type": "aix",
      "protocol": "fc",
      "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
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 <code>svm.name</code> or <code>svm.uuid</code> must be supplied.
5374901	Either <code>lun.uuid</code> or <code>lun.name</code> must be provided to create a LUN map.
5374902	Either <code>igroup.uuid</code> or <code>igroup.name</code> must be provided to create a LUN map.
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 <code>vvvol</code> cannot be mapped.
5374573	A node has no interface configured with the iSCSI or Fibre Channel protocols for the specified SVM.

Error Code	Description
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.
1254207	The LUN is already mapped to the same initiator group.

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	<a href="#">href</a>	

igroup

The initiator group to which the LUN is mapped. Required in POST by supplying either the `igroup.uuid`, `igroup.name`, or both.

Name	Type	Description
_links	<a href="#">_links</a>	
initiators	array[string]	The initiators that are members of the initiator group.
name	string	The name of the initiator group. Valid in POST.
os_type	string	The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system.
protocol	string	The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group.
uuid	string	The unique identifier of the initiator group. Valid in POST.

node

The LUN node.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name the LUN's node.

Name	Type	Description
uuid	string	The unique identifier of the LUN node.

lun

The LUN to which the initiator group is mapped. Required in POST by supplying either the `lun.uuid`, `lun.name`, or both.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	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	<a href="#">node</a>	The LUN node.
uuid	string	The unique identifier of the LUN. Valid in POST.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
igroup	<a href="#">igroup</a>	The initiator group to which the LUN is mapped. Required in POST by supplying either the <code>igroup.uuid</code> , <code>igroup.name</code> , or both.
logical_unit_number	integer	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. <ul style="list-style-type: none"> <li>• example: 1</li> <li>• max<b>Value</b>: 4095</li> <li>• min<b>Value</b>: 0</li> <li>• readCreate: 1</li> </ul>
lun	<a href="#">lun</a>	The LUN to which the initiator group is mapped. Required in POST by supplying either the <code>lun.uuid</code> , <code>lun.name</code> , or both.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete a LUN map

```
DELETE /protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}
```

Deletes a LUN map.

#### Related ONTAP commands

- [lun mapping delete](#)

#### Learn more

- [DOC /protocols/san/lun-maps](#)

#### Parameters

Name	Type	In	Required	Description
lun.uuid	string	path	True	
igroup.uuid	string	path	True	

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default
```

#### ONTAP Error Response Codes

Error Code	Description
5374875	The LUN was not found.

Name	Type	Description
error	<a href="#">error</a>	

### 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.



## Retrieve a LUN map

GET /protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}

Retrieves a LUN map.

### Related ONTAP commands

- `lun mapping show`

### Learn more

- [DOC /protocols/san/lun-maps](#)

### Parameters

Name	Type	In	Required	Description
lun.uuid	string	path	True	The unique identifier of the LUN.
igroup.uuid	string	path	True	The unique identifier of the igroup.
fields	array[string]	query	False	Specify the fields to return.

### Response

Status: 200, Ok

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>igroup</code>	<code>igroup</code>	The initiator group to which the LUN is mapped. Required in POST by supplying either the <code>igroup.uuid</code> , <code>igroup.name</code> , or both.

Name	Type	Description
logical_unit_number	integer	<p>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.</p> <ul style="list-style-type: none"> <li>• example: 1</li> <li>• maxValue: 4095</li> <li>• minValue: 0</li> <li>• readCreate: 1</li> </ul>
lun	lun	<p>The LUN to which the initiator group is mapped. Required in POST by supplying either the <code>lun.uuid</code>, <code>lun.name</code>, or both.</p>
svm	svm	<p>SVM, applies only to SVM-scoped objects.</p>

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "igroup": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "initiators": {
    },
    "name": "igroup1",
    "os_type": "aix",
    "protocol": "fc",
    "uuid": "1ad8544d-8cd1-91e0-9e1c-723478563412"
  },
  "logical_unit_number": 1,
  "lun": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "/vol/volume1/qtrees1/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.

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	<a href="#">href</a>	

igroup

The initiator group to which the LUN is mapped. Required in POST by supplying either the `igroup.uuid`, `igroup.name`, or both.

Name	Type	Description
_links	<a href="#">_links</a>	
initiators	array[string]	The initiators that are members of the initiator group.
name	string	The name of the initiator group. Valid in POST.
os_type	string	The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system.
protocol	string	The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group.
uuid	string	The unique identifier of the initiator group. Valid in POST.

node

The LUN node.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name the LUN's node.

Name	Type	Description
uuid	string	The unique identifier of the LUN node.

lun

The LUN to which the initiator group is mapped. Required in POST by supplying either the `lun.uuid`, `lun.name`, or both.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	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	<a href="#">node</a>	The LUN node.
uuid	string	The unique identifier of the LUN. Valid in POST.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage LUNs

### 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.

#### Examples

##### Creating a LUN

This example creates a 300 gigabyte, thin-provisioned LUN in SVM *svm1*, volume *vol1*, 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/vol1/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
  }
},
"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 to [/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 over-write an existing LUN as a clone of another. You do this as a PATCH on the LUN to overwrite — a PATCH to [/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, the patched LUN's data is over-written 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 patch 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 [/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 [/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'
```

## Retrieve LUNs

GET /storage/luns

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 [DOC Requesting specific fields](#) to learn more.

- `auto_delete`
- `lun_maps.*`
- `movement.*`
- `status.mapped`

### Related ONTAP commands

- `lun mapping show`
- `lun move show`
- `lun show`
- `volume file clone show-autodelete`

### Learn more

- [DOC /storage/luns](#)

### Parameters

Name	Type	In	Required	Description
<code>status.read_only</code>	boolean	query	False	Filter by <code>status.read_only</code>
<code>status.state</code>	string	query	False	Filter by <code>status.state</code>
<code>status.container_state</code>	string	query	False	Filter by <code>status.container_state</code>
<code>status.mapped</code>	boolean	query	False	Filter by <code>status.mapped</code>
<code>location.volume.name</code>	string	query	False	Filter by <code>location.volume.name</code>



<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
location.volume.uuid	string	query	False	Filter by location.volume.uuid
location.logical_unit	string	query	False	Filter by location.logical_unit
location.qtree.id	integer	query	False	Filter by location.qtree.id
location.qtree.name	string	query	False	Filter by location.qtree.name
qos_policy.name	string	query	False	Filter by qos_policy.name
qos_policy.uuid	string	query	False	Filter by qos_policy.uuid
serial_number	string	query	False	Filter by serial_number
auto_delete	boolean	query	False	Filter by auto_delete
os_type	string	query	False	Filter by os_type
enabled	boolean	query	False	Filter by enabled
movement.max_throughput	string	query	False	Filter by movement.max_throughput
movement.paths.source	string	query	False	Filter by movement.paths.source
movement.paths.destination	string	query	False	Filter by movement.paths.destination
movement.progress.failure.message	string	query	False	Filter by movement.progress.failure.message
movement.progress.failure.target	string	query	False	Filter by movement.progress.failure.target

<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
movement.progress.failure.arguments.message	string	query	False	Filter by movement.progress.failure.arguments.message
movement.progress.failure.arguments.code	string	query	False	Filter by movement.progress.failure.arguments.code
movement.progress.failure.code	string	query	False	Filter by movement.progress.failure.code
movement.progress.percent_complete	integer	query	False	Filter by movement.progress.percent_complete
movement.progress.volume_snapshot_blocked	boolean	query	False	Filter by movement.progress.volume_snapshot_blocked
movement.progress.elapsed	integer	query	False	Filter by movement.progress.elapsed
movement.progress.state	string	query	False	Filter by movement.progress.state
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
space.guarantee.reserved	boolean	query	False	Filter by space.guarantee.reserved
space.guarantee.requested	boolean	query	False	Filter by space.guarantee.requested
space.size	integer	query	False	Filter by space.size
space.used	integer	query	False	Filter by space.used

Name	Type	In	Required	Description
name	string	query	False	Filter by name
lun_maps.logical_unit_number	integer	query	False	Filter by lun_maps.logical_unit_number
lun_maps.igroup.name	string	query	False	Filter by lun_maps.igroup.name
lun_maps.igroup.uuid	string	query	False	Filter by lun_maps.igroup.uuid
comment	string	query	False	Filter by comment
uuid	string	query	False	Filter by uuid
class	string	query	False	Filter by class
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.
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.

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	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">lun</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "class": "regular",
    "clone": {
      "source": {
        "name": "/vol/volume1/lun1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "comment": "string",
    "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
},
"movement": {
  "paths": {
    "destination": "/vol/vol1/lun1",
    "source": "/vol/vol2/lun2"
  },
  "progress": {
    "elapsed": 0,
    "failure": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    },
    "percent_complete": 0,
    "state": "preparing"
  }
},
"name": "/vol/volume1/qtree1/lun1",
"os_type": "aix",
"qos_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "qos1",

```

```
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "serial_number": "string",
  "space": {
    "size": 1073741824,
    "used": 0
  },
  "status": {
    "container_state": "online",
    "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

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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.

Name	Type	Description
name	string	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	The unique identifier of the clone source LUN. Valid in POST and PATCH.

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.

Name	Type	Description
source	<a href="#">source</a>	<p>The source LUN for a LUN clone operation. This can be specified using property <code>clone.source.uuid</code> or <code>clone.source.name</code>. If both properties are supplied, they must refer to the same LUN.</p> <p>Valid in POST to create a new LUN as a clone of the source.</p> <p>Valid in PATCH to overwrite an existing LUN's data as a clone of another.</p>

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree.

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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

location

The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

Name	Type	Description
logical_unit	string	<p>The base name component of the LUN. Valid in POST and PATCH.</p> <p>If properties <code>name</code> and <code>location.logical_unit</code> are specified in the same request, they must refer to the base name.</p> <p>A PATCH that modifies the base name of the LUN is considered a rename operation.</p>
qtree	qtree	<p>The qtree in which the LUN is optionally located. Valid in POST and PATCH.</p> <p>If properties <code>name</code> and <code>location.qtree.name</code> and/or <code>location.qtree.uuid</code> are specified in the same request, they must refer to the same qtree.</p> <p>A PATCH that modifies the qtree of the LUN is considered a rename operation.</p>

Name	Type	Description
volume	<a href="#">volume</a>	<p>The volume in which the LUN is located. Valid in POST and PATCH.</p> <p>If properties <code>name</code> and <code>location.volume.name</code> and/or <code>location.volume.uuid</code> are specified in the same request, they must refer to the same volume.</p> <p>A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.</p>

#### igroup

The initiator group to which the LUN is mapped.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the initiator group.
uuid	string	The unique identifier of the initiator group.

#### lun\_maps

A LUN map with which the LUN is associated.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
igroup	<a href="#">igroup</a>	The initiator group to which the LUN is mapped.
logical_unit_number	integer	The logical unit number assigned to the LUN for initiators in the initiator group.

#### paths

The fully qualified LUN path names involved in the LUN movement.

Name	Type	Description
destination	string	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	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.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### failure

Error information provided if the asynchronous LUN movement operation fails.

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

#### progress

Name	Type	Description
elapsed	integer	The amount of time, in seconds, that has elapsed since the start of the LUN movement.

Name	Type	Description
failure	<a href="#">failure</a>	Error information provided if the asynchronous LUN movement operation fails.
percent_complete	integer	The percentage complete of the LUN movement.
state	string	The state of the LUN movement.  Valid in PATCH when an LUN movement is active. Set to <i>paused</i> to pause a LUN movement. Set to <i>replicating</i> to resume a paused LUN movement.
volume_snapshot_blocked	boolean	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

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 [DOC Requesting specific fields](#) to learn more.

Name	Type	Description
max_throughput	string	<p>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.</p> <p>If this property is not specified in a POST that begins a LUN movement, throttling is not applied to the data transfer.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <p>This property is valid only in a POST that begins a LUN movement or a PATCH when a LUN movement is already in process.</p>
paths	<a href="#">paths</a>	The fully qualified LUN path names involved in the LUN movement.
progress	<a href="#">progress</a>	

#### 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.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	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	The unique identifier of the QoS policy. Valid in POST and PATCH.

### guarantee

Properties that request and report the space guarantee for the LUN.

Name	Type	Description
requested	boolean	The requested space reservation policy for the LUN. If <i>true</i> , a space reservation is requested for the LUN; if <i>false</i> , 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	Reports if the LUN is space guaranteed.  If <i>true</i> , a space guarantee is requested and the containing volume and aggregate support the request. If <i>false</i> , 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

The storage space related properties of the LUN.



Name	Type	Description
guarantee	<a href="#">guarantee</a>	Properties that request and report the space guarantee for the LUN.
size	integer	<p>The total provisioned size of the LUN. The LUN size can be increased but not be made smaller using the REST interface.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• example: 1073741824</li> </ul>
used	integer	<p>The amount of space consumed by the main data stream of the LUN.</p> <p>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.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

status

Status information about the LUN.

Name	Type	Description
container_state	string	The state of the volume and aggregate that contain the LUN. LUNs are only available when their containers are available.
mapped	boolean	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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.
read_only	boolean	Reports if the LUN allows only read access.
state	string	The state of the LUN. Normal states for a LUN are <i>online</i> and <i>offline</i> . Other states indicate errors.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
auto_delete	boolean	<p>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.</p> <p>When set to <i>true</i>, 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.</p> <p>This property is optional in POST and PATCH. The default value for a new LUN is <i>false</i>.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
class	string	The class of LUN. Only <i>regular</i> LUNs can be created using the REST API.

Name	Type	Description
clone	clone	<p>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: <code>auto_delete</code>, <code>qos_policy</code>, and <code>space.guarantee.requested</code>.</p> <p>When used in a PATCH, the patched LUN's data is overwritten as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: <code>class</code>, <code>auto_delete</code>, <code>lun_maps</code>, <code>serial_number</code>, <code>status.state</code>, and <code>uuid</code>.</p> <p>Persistent reservations for the patched LUN are also preserved.</p>
comment	string	A configurable comment available for use by the administrator. Valid in POST and PATCH.
enabled	boolean	<p>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 <code>state</code> property to determine if the LUN is administratively disabled (<i>offline</i>) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the <code>enabled</code> property to <i>true</i> or brought administratively offline by setting the <code>enabled</code> property to <i>false</i>. Upon creation, a LUN is enabled by default. Valid in PATCH.</p>

Name	Type	Description
location	<a href="#">location</a>	The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.
lun_maps	array[ <a href="#">lun_maps</a> ]	<p>The LUN maps with which the LUN is associated.</p> <p>There is an added cost to retrieving property values for <code>lun_maps</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
movement	movement	<p>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 <code>name</code>, <code>location.volume.uuid</code>, or <code>location.volume.name</code>. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.</p> <p>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 <code>movement</code> sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.</p> <p>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 <code>movement</code> properties. The LUN movement operation can be further modified using a PATCH on the properties on the <code>movement</code> sub-object.</p> <p>There is added cost to retrieving property values for <code>movement</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
name	string	<p>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.</p> <p>A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.</p> <p>A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation.</p>
os_type	string	<p>The operating system type of the LUN.</p> <p>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone.</p>
qos_policy	qos_policy	<p>The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property <code>qos_policy.uuid</code> and <code>qos_policy.name</code> 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 <code>qos_policy.name</code> to an empty string ("") in a PATCH request. Valid in POST and PATCH.</p> <p>Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.</p>

Name	Type	Description
serial_number	string	The LUN serial number. The serial number is generated by ONTAP when the LUN is created. <ul style="list-style-type: none"> <li>• maxLength: 12</li> <li>• minLength: 12</li> <li>• readOnly: 1</li> </ul>
space	<a href="#">space</a>	The storage space related properties of the LUN.
status	<a href="#">status</a>	Status information about the LUN.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created. <ul style="list-style-type: none"> <li>• example: 1cd8a442-86d1-11e0-ae1c-123478563412</li> <li>• readOnly: 1</li> </ul>

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create a LUN

POST /storage/luns

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

- [DOC /storage/luns](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
auto_delete	boolean	<p>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.</p> <p>When set to <i>true</i>, 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.</p> <p>This property is optional in POST and PATCH. The default value for a new LUN is <i>false</i>.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
class	string	<p>The class of LUN. Only <i>regular</i> LUNs can be created using the REST API.</p>

Name	Type	Description
clone	clone	<p>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: <code>auto_delete</code>, <code>qos_policy</code>, and <code>space.guarantee.requested</code>.</p> <p>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: <code>class</code>, <code>auto_delete</code>, <code>lun_maps</code>, <code>serial_number</code>, <code>status.state</code>, and <code>uuid</code>.</p> <p>Persistent reservations for the patched LUN are also preserved.</p>
comment	string	A configurable comment available for use by the administrator. Valid in POST and PATCH.
enabled	boolean	<p>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 <code>state</code> property to determine if the LUN is administratively disabled (<i>offline</i>) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the <code>enabled</code> property to <i>true</i> or brought administratively offline by setting the <code>enabled</code> property to <i>false</i>. Upon creation, a LUN is enabled by default. Valid in PATCH.</p>
location	location	The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

Name	Type	Description
lun_maps	array[lun_maps]	<p>The LUN maps with which the LUN is associated.</p> <p>There is an added cost to retrieving property values for <code>lun_maps</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
movement	movement	<p>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 <code>name</code>, <code>location.volume.uuid</code>, or <code>location.volume.name</code>. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.</p> <p>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 <code>movement</code> sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.</p> <p>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 <code>movement</code> properties. The LUN movement operation can be further modified using a PATCH on the properties on the <code>movement</code> sub-object.</p> <p>There is added cost to retrieving property values for <code>movement</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
name	string	<p>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.</p> <p>A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.</p> <p>A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation.</p>
os_type	string	<p>The operating system type of the LUN.</p> <p>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone.</p>
qos_policy	<a href="#">qos_policy</a>	<p>The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property <code>qos_policy.uuid</code> and <code>qos_policy.name</code> 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 <code>qos_policy.name</code> to an empty string ("") in a PATCH request. Valid in POST and PATCH.</p> <p>Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.</p>
serial_number	string	<p>The LUN serial number. The serial number is generated by ONTAP when the LUN is created.</p> <ul style="list-style-type: none"> <li>• maxLength: 12</li> <li>• minLength: 12</li> <li>• readOnly: 1</li> </ul>

Name	Type	Description
space	space	The storage space related properties of the LUN.
status	status	Status information about the LUN.
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	<p>The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.</p> <ul style="list-style-type: none"> <li>• example: 1cd8a442-86d1-11e0-ae1c-123478563412</li> <li>• readOnly: 1</li> </ul>

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "class": "regular",
  "clone": {
    "source": {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "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
},
"movement": {
  "paths": {
    "destination": "/vol/vol1/lun1",
    "source": "/vol/vol2/lun2"
  },
  "progress": {
    "elapsed": 0,
    "failure": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    },
    "percent_complete": 0,
    "state": "preparing"
  }
},
"name": "/vol/volume1/qtree1/lun1",
"os_type": "aix",
"qos_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "qos1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"serial_number": "string",
"space": {
  "size": 1073741824,
  "used": 0
},
"status": {
  "container_state": "online",

```

```
    "state": "online"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

### Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records.
records	array[ <a href="#">lun</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "class": "regular",
    "clone": {
      "source": {
        "name": "/vol/volume1/lun1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "comment": "string",
    "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
},
"movement": {
  "paths": {
    "destination": "/vol/vol1/lun1",
    "source": "/vol/vol2/lun2"
  },
  "progress": {
    "elapsed": 0,
    "failure": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    },
    "percent_complete": 0,
    "state": "preparing"
  }
},
"name": "/vol/volume1/mtree1/lun1",
"os_type": "aix",
"qos_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "qos1",

```

```

    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "serial_number": "string",
  "space": {
    "size": 1073741824,
    "used": 0
  },
  "status": {
    "container_state": "online",
    "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
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.
5374863	An error occurred after successfully creating the LUN. Some properties were not set.
5374886	An error occurred after successfully creating the LUN preventing the retrieval of its properties.
5374874	The specified <code>clone.source.uuid</code> and <code>clone.source.name</code> do not refer to the same LUN.

Error Code	Description
5374875	The specified <code>clone.source</code> was not found.
5374876	The specified <code>clone.source</code> was not found.
917927	The specified volume was not found.
918236	The specified <code>location.volume.uuid</code> and <code>location.volume.name</code> do not refer to the same volume.
5374858	The volume specified by <code>name</code> is not the same as that specified by <code>location.volume</code> .
5242927	The specified <code>qtree</code> was not found.
5242950	The specified <code>location.qtree.id</code> and <code>location.qtree.name</code> do not refer to the same <code>qtree</code> .
5374860	The <code>qtree</code> specified by <code>name</code> is not the same as that specified by <code>location.qtree</code> .
5374861	The LUN base name specified by <code>name</code> is not the same as that specified by <code>location.logical_unit</code> .
13565952	The LUN clone request failed.
5374130	An invalid size value was provided.
5374241	A size value with invalid units was provided.
5374125	The specified size is too large for the LUN.
5374124	The specified size is too small for the LUN.
5374242	A LUN or NVMe namespace already exists at the specified path.
5374707	Creating a LUN in the specific volume is not allowed because the volume is reserved for an application.
5374883	The property cannot be specified when creating a LUN clone. The <code>target</code> property of the error object identifies the property.
5374859	No volume was specified for the LUN.
5374884	The property is required except when creating a LUN clone. The <code>target</code> property of the error object identifies the property.
5374862	No LUN path base name was provided for the LUN.
5374123	A negative size was provided for the LUN.
5374352	An invalid name was provided for the LUN.
5374129	LUNs cannot be created on a load sharing mirror volume.

Error Code	Description
5374237	LUNs cannot be created on an SVM root volume.
5374121	A LUN name can only contain characters A-Z, a-z, 0-9, "-", ".", "_", "{" and "}".
5374238	LUNs cannot be created in Snapshot copies.
5374899	The <code>clone.source.uuid</code> property is not supported when specifying a source LUN from a Snapshot copy.

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	<a href="#">href</a>	

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.

Name	Type	Description
name	string	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	The unique identifier of the clone source LUN. Valid in POST and PATCH.

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.



Name	Type	Description
source	<a href="#">source</a>	<p>The source LUN for a LUN clone operation. This can be specified using property <code>clone.source.uuid</code> or <code>clone.source.name</code>. If both properties are supplied, they must refer to the same LUN.</p> <p>Valid in POST to create a new LUN as a clone of the source.</p> <p>Valid in PATCH to overwrite an existing LUN's data as a clone of another.</p>

### 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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree.

### 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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

location

The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

Name	Type	Description
logical_unit	string	<p>The base name component of the LUN. Valid in POST and PATCH.</p> <p>If properties <code>name</code> and <code>location.logical_unit</code> are specified in the same request, they must refer to the base name.</p> <p>A PATCH that modifies the base name of the LUN is considered a rename operation.</p>
qtree	qtree	<p>The qtree in which the LUN is optionally located. Valid in POST and PATCH.</p> <p>If properties <code>name</code> and <code>location.qtree.name</code> and/or <code>location.qtree.uuid</code> are specified in the same request, they must refer to the same qtree.</p> <p>A PATCH that modifies the qtree of the LUN is considered a rename operation.</p>

Name	Type	Description
volume	<a href="#">volume</a>	<p>The volume in which the LUN is located. Valid in POST and PATCH.</p> <p>If properties <code>name</code> and <code>location.volume.name</code> and/or <code>location.volume.uuid</code> are specified in the same request, they must refer to the same volume.</p> <p>A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.</p>

#### igroup

The initiator group to which the LUN is mapped.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the initiator group.
uuid	string	The unique identifier of the initiator group.

#### lun\_maps

A LUN map with which the LUN is associated.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
igroup	<a href="#">igroup</a>	The initiator group to which the LUN is mapped.
logical_unit_number	integer	The logical unit number assigned to the LUN for initiators in the initiator group.

#### paths

The fully qualified LUN path names involved in the LUN movement.

Name	Type	Description
destination	string	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	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.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### failure

Error information provided if the asynchronous LUN movement operation fails.

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

#### progress

Name	Type	Description
elapsed	integer	The amount of time, in seconds, that has elapsed since the start of the LUN movement.

Name	Type	Description
failure	<a href="#">failure</a>	Error information provided if the asynchronous LUN movement operation fails.
percent_complete	integer	The percentage complete of the LUN movement.
state	string	The state of the LUN movement.  Valid in PATCH when an LUN movement is active. Set to <i>paused</i> to pause a LUN movement. Set to <i>replicating</i> to resume a paused LUN movement.
volume_snapshot_blocked	boolean	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

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 [DOC Requesting specific fields](#) to learn more.

Name	Type	Description
max_throughput	string	<p>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.</p> <p>If this property is not specified in a POST that begins a LUN movement, throttling is not applied to the data transfer.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <p>This property is valid only in a POST that begins a LUN movement or a PATCH when a LUN movement is already in process.</p>
paths	<a href="#">paths</a>	The fully qualified LUN path names involved in the LUN movement.
progress	<a href="#">progress</a>	

#### 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.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	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	The unique identifier of the QoS policy. Valid in POST and PATCH.

### guarantee

Properties that request and report the space guarantee for the LUN.

Name	Type	Description
requested	boolean	The requested space reservation policy for the LUN. If <i>true</i> , a space reservation is requested for the LUN; if <i>false</i> , 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	Reports if the LUN is space guaranteed.  If <i>true</i> , a space guarantee is requested and the containing volume and aggregate support the request. If <i>false</i> , 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

The storage space related properties of the LUN.

Name	Type	Description
guarantee	guarantee	Properties that request and report the space guarantee for the LUN.
size	integer	<p>The total provisioned size of the LUN. The LUN size can be increased but not be made smaller using the REST interface.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• example: 1073741824</li> </ul>
used	integer	<p>The amount of space consumed by the main data stream of the LUN.</p> <p>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.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

status

Status information about the LUN.



Name	Type	Description
container_state	string	The state of the volume and aggregate that contain the LUN. LUNs are only available when their containers are available.
mapped	boolean	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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.
read_only	boolean	Reports if the LUN allows only read access.
state	string	The state of the LUN. Normal states for a LUN are <i>online</i> and <i>offline</i> . Other states indicate errors.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
auto_delete	boolean	<p>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.</p> <p>When set to <i>true</i>, 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.</p> <p>This property is optional in POST and PATCH. The default value for a new LUN is <i>false</i>.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
class	string	The class of LUN. Only <i>regular</i> LUNs can be created using the REST API.

Name	Type	Description
clone	clone	<p>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: <code>auto_delete</code>, <code>qos_policy</code>, and <code>space.guarantee.requested</code>.</p> <p>When used in a PATCH, the patched LUN's data is overwritten as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: <code>class</code>, <code>auto_delete</code>, <code>lun_maps</code>, <code>serial_number</code>, <code>status.state</code>, and <code>uuid</code>.</p> <p>Persistent reservations for the patched LUN are also preserved.</p>
comment	string	A configurable comment available for use by the administrator. Valid in POST and PATCH.
enabled	boolean	<p>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 <code>state</code> property to determine if the LUN is administratively disabled (<i>offline</i>) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the <code>enabled</code> property to <i>true</i> or brought administratively offline by setting the <code>enabled</code> property to <i>false</i>. Upon creation, a LUN is enabled by default. Valid in PATCH.</p>

Name	Type	Description
location	<a href="#">location</a>	The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.
lun_maps	array[ <a href="#">lun_maps</a> ]	<p>The LUN maps with which the LUN is associated.</p> <p>There is an added cost to retrieving property values for <code>lun_maps</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
movement	movement	<p>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 <code>name</code>, <code>location.volume.uuid</code>, or <code>location.volume.name</code>. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.</p> <p>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 <code>movement</code> sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.</p> <p>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 <code>movement</code> properties. The LUN movement operation can be further modified using a PATCH on the properties on the <code>movement</code> sub-object.</p> <p>There is added cost to retrieving property values for <code>movement</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
name	string	<p>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.</p> <p>A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.</p> <p>A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation.</p>
os_type	string	<p>The operating system type of the LUN.</p> <p>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone.</p>
qos_policy	qos_policy	<p>The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property <code>qos_policy.uuid</code> and <code>qos_policy.name</code> 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 <code>qos_policy.name</code> to an empty string ("") in a PATCH request. Valid in POST and PATCH.</p> <p>Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.</p>

Name	Type	Description
serial_number	string	The LUN serial number. The serial number is generated by ONTAP when the LUN is created. <ul style="list-style-type: none"> <li>• maxLength: 12</li> <li>• minLength: 12</li> <li>• readOnly: 1</li> </ul>
space	<a href="#">space</a>	The storage space related properties of the LUN.
status	<a href="#">status</a>	Status information about the LUN.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created. <ul style="list-style-type: none"> <li>• example: 1cd8a442-86d1-11e0-ae1c-123478563412</li> <li>• readOnly: 1</li> </ul>

#### [\\_links](#)

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete a LUN

```
DELETE /storage/luns/{uuid}
```

Deletes a LUN.

### Related ONTAP commands

- `lun delete`

### Learn more

- [DOC /storage/luns](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier of the LUN.
allow_delete_while_mapped	boolean	query	False	Allow 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. <b>This parameter should be used with caution.</b> <ul style="list-style-type: none"><li>• Default value:</li></ul>

### Response

```
Status: 200, Ok
```

### Error

```
Status: Default
```

### ONTAP Error Response Codes



Error Code	Description
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.
5374705	Deleting the LUN is not allowed because it is part of an application.
1254197	The LUN is mapped and cannot be deleted without specifying the <code>allow_delete_while_mapped</code> query parameter.

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve LUN properties or data

GET /storage/luns/{uuid}

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 [DOC Requesting specific fields](#) to learn more.

- `auto_delete`
- `lun_maps.*`
- `movement.*`
- `status.mapped`

### Related ONTAP commands

- `lun mapping show`
- `lun move show`
- `lun show`
- `volume file clone show-autodelete`

## Learn more

- [DOC /storage/luns](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier of the LUN to retrieve.
fields	array[string]	query	False	Specify the fields to return.

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
auto_delete	boolean	<p>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.</p> <p>When set to <i>true</i>, 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.</p> <p>This property is optional in POST and PATCH. The default value for a new LUN is <i>false</i>.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
class	string	The class of LUN. Only <i>regular</i> LUNs can be created using the REST API.
clone	clone	<p>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: <code>auto_delete</code>, <code>qos_policy</code>, and <code>space.guarantee.requested</code>.</p> <p>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: <code>class</code>, <code>auto_delete</code>, <code>lun_maps</code>, <code>serial_number</code>, <code>status.state</code>, and <code>uuid</code>.</p> <p>Persistent reservations for the patched LUN are also preserved.</p>
comment	string	A configurable comment available for use by the administrator. Valid in POST and PATCH.
enabled	boolean	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 <code>state</code> property to determine if the LUN is administratively disabled ( <i>offline</i> ) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the <code>enabled</code> property to <i>true</i> or brought administratively offline by setting the <code>enabled</code> property to <i>false</i> . Upon creation, a LUN is enabled by default. Valid in PATCH.

Name	Type	Description
location	<a href="#">location</a>	The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.
lun_maps	array[ <a href="#">lun_maps</a> ]	<p>The LUN maps with which the LUN is associated.</p> <p>There is an added cost to retrieving property values for <code>lun_maps</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
movement	movement	<p>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 <code>name</code>, <code>location.volume.uuid</code>, or <code>location.volume.name</code>. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.</p> <p>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 <code>movement</code> sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.</p> <p>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 <code>movement</code> properties. The LUN movement operation can be further modified using a PATCH on the properties on the <code>movement</code> sub-object.</p> <p>There is added cost to retrieving property values for <code>movement</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
name	string	<p>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.</p> <p>A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.</p> <p>A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation.</p>
os_type	string	<p>The operating system type of the LUN.</p> <p>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone.</p>
qos_policy	<a href="#">qos_policy</a>	<p>The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property <code>qos_policy.uuid</code> and <code>qos_policy.name</code> 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 <code>qos_policy.name</code> to an empty string ("") in a PATCH request. Valid in POST and PATCH.</p> <p>Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.</p>
serial_number	string	<p>The LUN serial number. The serial number is generated by ONTAP when the LUN is created.</p> <ul style="list-style-type: none"> <li>• maxLength: 12</li> <li>• minLength: 12</li> <li>• readOnly: 1</li> </ul>

Name	Type	Description
space	space	The storage space related properties of the LUN.
status	status	Status information about the LUN.
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	<p>The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.</p> <ul style="list-style-type: none"> <li>• example: 1cd8a442-86d1-11e0-ae1c-123478563412</li> <li>• readOnly: 1</li> </ul>



## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "class": "regular",
  "clone": {
    "source": {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "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
},
"movement": {
  "paths": {
    "destination": "/vol/vol1/lun1",
    "source": "/vol/vol2/lun2"
  },
  "progress": {
    "elapsed": 0,
    "failure": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    },
    "percent_complete": 0,
    "state": "preparing"
  }
},
"name": "/vol/volume1/qtree1/lun1",
"os_type": "aix",
"qos_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "qos1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"serial_number": "string",
"space": {
  "size": 1073741824,
  "used": 0
},
"status": {
  "container_state": "online",

```

```

    "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.

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	<a href="#">href</a>	

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.

Name	Type	Description
name	string	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	The unique identifier of the clone source LUN. Valid in POST and PATCH.

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.

Name	Type	Description
source	<a href="#">source</a>	<p>The source LUN for a LUN clone operation. This can be specified using property <code>clone.source.uuid</code> or <code>clone.source.name</code>. If both properties are supplied, they must refer to the same LUN.</p> <p>Valid in POST to create a new LUN as a clone of the source.</p> <p>Valid in PATCH to overwrite an existing LUN's data as a clone of another.</p>

### 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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree.

### 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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

location

The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

Name	Type	Description
logical_unit	string	<p>The base name component of the LUN. Valid in POST and PATCH.</p> <p>If properties <code>name</code> and <code>location.logical_unit</code> are specified in the same request, they must refer to the base name.</p> <p>A PATCH that modifies the base name of the LUN is considered a rename operation.</p>
qtree	qtree	<p>The qtree in which the LUN is optionally located. Valid in POST and PATCH.</p> <p>If properties <code>name</code> and <code>location.qtree.name</code> and/or <code>location.qtree.uuid</code> are specified in the same request, they must refer to the same qtree.</p> <p>A PATCH that modifies the qtree of the LUN is considered a rename operation.</p>

Name	Type	Description
volume	<a href="#">volume</a>	<p>The volume in which the LUN is located. Valid in POST and PATCH.</p> <p>If properties <code>name</code> and <code>location.volume.name</code> and/or <code>location.volume.uuid</code> are specified in the same request, they must refer to the same volume.</p> <p>A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.</p>

#### igroup

The initiator group to which the LUN is mapped.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the initiator group.
uuid	string	The unique identifier of the initiator group.

#### lun\_maps

A LUN map with which the LUN is associated.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
igroup	<a href="#">igroup</a>	The initiator group to which the LUN is mapped.
logical_unit_number	integer	The logical unit number assigned to the LUN for initiators in the initiator group.

#### paths

The fully qualified LUN path names involved in the LUN movement.



Name	Type	Description
destination	string	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	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.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### failure

Error information provided if the asynchronous LUN movement operation fails.

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

#### progress

Name	Type	Description
elapsed	integer	The amount of time, in seconds, that has elapsed since the start of the LUN movement.

Name	Type	Description
failure	<a href="#">failure</a>	Error information provided if the asynchronous LUN movement operation fails.
percent_complete	integer	The percentage complete of the LUN movement.
state	string	The state of the LUN movement.  Valid in PATCH when an LUN movement is active. Set to <i>paused</i> to pause a LUN movement. Set to <i>replicating</i> to resume a paused LUN movement.
volume_snapshot_blocked	boolean	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

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 [DOC Requesting specific fields](#) to learn more.

Name	Type	Description
max_throughput	string	<p>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.</p> <p>If this property is not specified in a POST that begins a LUN movement, throttling is not applied to the data transfer.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <p>This property is valid only in a POST that begins a LUN movement or a PATCH when a LUN movement is already in process.</p>
paths	<a href="#">paths</a>	The fully qualified LUN path names involved in the LUN movement.
progress	<a href="#">progress</a>	

#### 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.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	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	The unique identifier of the QoS policy. Valid in POST and PATCH.

### guarantee

Properties that request and report the space guarantee for the LUN.

Name	Type	Description
requested	boolean	The requested space reservation policy for the LUN. If <i>true</i> , a space reservation is requested for the LUN; if <i>false</i> , 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	Reports if the LUN is space guaranteed.  If <i>true</i> , a space guarantee is requested and the containing volume and aggregate support the request. If <i>false</i> , 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

The storage space related properties of the LUN.

Name	Type	Description
guarantee	<a href="#">guarantee</a>	Properties that request and report the space guarantee for the LUN.
size	integer	<p>The total provisioned size of the LUN. The LUN size can be increased but not be made smaller using the REST interface.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• example: 1073741824</li> </ul>
used	integer	<p>The amount of space consumed by the main data stream of the LUN.</p> <p>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.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

status

Status information about the LUN.

Name	Type	Description
container_state	string	The state of the volume and aggregate that contain the LUN. LUNs are only available when their containers are available.
mapped	boolean	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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.
read_only	boolean	Reports if the LUN allows only read access.
state	string	The state of the LUN. Normal states for a LUN are <i>online</i> and <i>offline</i> . Other states indicate errors.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

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

## Update an existing LUN

PATCH /storage/luns/{uuid}

Updates the properties of a LUN. PATCH 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

- [DOC /storage/luns](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	The unique identifier of the LUN to update.

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
auto_delete	boolean	<p>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.</p> <p>When set to <i>true</i>, 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.</p> <p>This property is optional in POST and PATCH. The default value for a new LUN is <i>false</i>.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
class	string	<p>The class of LUN. Only <i>regular</i> LUNs can be created using the REST API.</p>



Name	Type	Description
clone	clone	<p>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: <code>auto_delete</code>, <code>qos_policy</code>, and <code>space.guarantee.requested</code>.</p> <p>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: <code>class</code>, <code>auto_delete</code>, <code>lun_maps</code>, <code>serial_number</code>, <code>status.state</code>, and <code>uuid</code>.</p> <p>Persistent reservations for the patched LUN are also preserved.</p>
comment	string	A configurable comment available for use by the administrator. Valid in POST and PATCH.
enabled	boolean	<p>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 <code>state</code> property to determine if the LUN is administratively disabled (<i>offline</i>) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the <code>enabled</code> property to <i>true</i> or brought administratively offline by setting the <code>enabled</code> property to <i>false</i>. Upon creation, a LUN is enabled by default. Valid in PATCH.</p>
location	location	The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

Name	Type	Description
lun_maps	array[lun_maps]	<p>The LUN maps with which the LUN is associated.</p> <p>There is an added cost to retrieving property values for <code>lun_maps</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
movement	movement	<p>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 <code>name</code>, <code>location.volume.uuid</code>, or <code>location.volume.name</code>. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.</p> <p>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 <code>movement</code> sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.</p> <p>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 <code>movement</code> properties. The LUN movement operation can be further modified using a PATCH on the properties on the <code>movement</code> sub-object.</p> <p>There is added cost to retrieving property values for <code>movement</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
name	string	<p>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.</p> <p>A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.</p> <p>A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation.</p>
os_type	string	<p>The operating system type of the LUN.</p> <p>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone.</p>
qos_policy	<a href="#">qos_policy</a>	<p>The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property <code>qos_policy.uuid</code> and <code>qos_policy.name</code> 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 <code>qos_policy.name</code> to an empty string ("") in a PATCH request. Valid in POST and PATCH.</p> <p>Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.</p>
serial_number	string	<p>The LUN serial number. The serial number is generated by ONTAP when the LUN is created.</p> <ul style="list-style-type: none"> <li>• maxLength: 12</li> <li>• minLength: 12</li> <li>• readOnly: 1</li> </ul>

Name	Type	Description
space	space	The storage space related properties of the LUN.
status	status	Status information about the LUN.
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	<p>The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.</p> <ul style="list-style-type: none"> <li>• example: 1cd8a442-86d1-11e0-ae1c-123478563412</li> <li>• readOnly: 1</li> </ul>

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "class": "regular",
  "clone": {
    "source": {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "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
},
"movement": {
    "paths": {
        "destination": "/vol/vol1/lun1",
        "source": "/vol/vol2/lun2"
    },
    "progress": {
        "elapsed": 0,
        "failure": {
            "arguments": {
                "code": "string",
                "message": "string"
            },
            "code": "4",
            "message": "entry doesn't exist",
            "target": "uuid"
        },
        "percent_complete": 0,
        "state": "preparing"
    }
},
"name": "/vol/volume1/qtree1/lun1",
"os_type": "aix",
"qos_policy": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "qos1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"serial_number": "string",
"space": {
    "size": 1073741824,
    "used": 0
},
"status": {
    "container_state": "online",

```

```

    "state": "online"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resource/link"
      }
    },
    "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
5374889	An invalid value was specified for <code>movement.progress.state</code> . Active LUN movement operations can be PATCHed to only <i>paused</i> or <i>replicating</i> .
5374864	An error occurred after successfully overwriting data for the LUN as a clone. Some properties were not modified.
5374885	An error occurred after successfully modifying some of the properties of the LUN. 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 <code>clone.source.uuid</code> and <code>clone.source.name</code> do not refer to the same LUN.



Error Code	Description
5374875	The specified LUN was not found. This can apply to <code>clone.source</code> or the target LUN. The <code>target</code> property of the error object identifies the property.
5374876	The specified LUN was not found. This can apply to <code>clone.source</code> or the target LUN. The <code>target</code> property of the error object identifies the property.
917927	The specified volume was not found.
918236	The specified <code>location.volume.uuid</code> and <code>location.volume.name</code> do not refer to the same volume.
5374858	The volume specified by <code>name</code> is not the same as that specified by <code>location.volume</code> .
5242927	The specified <code>qtree</code> was not found.
5242950	The specified <code>location.qtree.id</code> and <code>location.qtree.name</code> do not refer to the same <code>qtree</code> .
5374860	The <code>qtree</code> specified by <code>name</code> is not the same as that specified by <code>location.qtree</code> .
5374861	The LUN base name specified by <code>name</code> is not the same as that specified by <code>location.logical_unit</code> .
13565952	The LUN clone request failed.
5374130	An invalid size value was provided.
5374241	A size value with invalid units was provided.
5374125	The specified LUN size is too large.
5374124	The specified LUN size is too small.
5374892	An attempt was made to reduce the size of a LUN.
5374480	Modifying the LUN is not allowed because it is in a foreign LUN import relationship.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

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.

Name	Type	Description
name	string	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	The unique identifier of the clone source LUN. Valid in POST and PATCH.

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.

Name	Type	Description
source	<a href="#">source</a>	<p>The source LUN for a LUN clone operation. This can be specified using property <code>clone.source.uuid</code> or <code>clone.source.name</code>. If both properties are supplied, they must refer to the same LUN.</p> <p>Valid in POST to create a new LUN as a clone of the source.</p> <p>Valid in PATCH to overwrite an existing LUN's data as a clone of another.</p>

### 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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree.

### 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.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

location

The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

Name	Type	Description
logical_unit	string	<p>The base name component of the LUN. Valid in POST and PATCH.</p> <p>If properties <code>name</code> and <code>location.logical_unit</code> are specified in the same request, they must refer to the base name.</p> <p>A PATCH that modifies the base name of the LUN is considered a rename operation.</p>
qtree	qtree	<p>The qtree in which the LUN is optionally located. Valid in POST and PATCH.</p> <p>If properties <code>name</code> and <code>location.qtree.name</code> and/or <code>location.qtree.uuid</code> are specified in the same request, they must refer to the same qtree.</p> <p>A PATCH that modifies the qtree of the LUN is considered a rename operation.</p>

Name	Type	Description
volume	<a href="#">volume</a>	<p>The volume in which the LUN is located. Valid in POST and PATCH.</p> <p>If properties <code>name</code> and <code>location.volume.name</code> and/or <code>location.volume.uuid</code> are specified in the same request, they must refer to the same volume.</p> <p>A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.</p>

#### igroup

The initiator group to which the LUN is mapped.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the initiator group.
uuid	string	The unique identifier of the initiator group.

#### lun\_maps

A LUN map with which the LUN is associated.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
igroup	<a href="#">igroup</a>	The initiator group to which the LUN is mapped.
logical_unit_number	integer	The logical unit number assigned to the LUN for initiators in the initiator group.

#### paths

The fully qualified LUN path names involved in the LUN movement.

Name	Type	Description
destination	string	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	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.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### failure

Error information provided if the asynchronous LUN movement operation fails.

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

#### progress

Name	Type	Description
elapsed	integer	The amount of time, in seconds, that has elapsed since the start of the LUN movement.

Name	Type	Description
failure	<a href="#">failure</a>	Error information provided if the asynchronous LUN movement operation fails.
percent_complete	integer	The percentage complete of the LUN movement.
state	string	The state of the LUN movement.  Valid in PATCH when an LUN movement is active. Set to <i>paused</i> to pause a LUN movement. Set to <i>replicating</i> to resume a paused LUN movement.
volume_snapshot_blocked	boolean	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

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 [DOC Requesting specific fields](#) to learn more.



Name	Type	Description
max_throughput	string	<p>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.</p> <p>If this property is not specified in a POST that begins a LUN movement, throttling is not applied to the data transfer.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <p>This property is valid only in a POST that begins a LUN movement or a PATCH when a LUN movement is already in process.</p>
paths	<a href="#">paths</a>	The fully qualified LUN path names involved in the LUN movement.
progress	<a href="#">progress</a>	

#### 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.

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	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	The unique identifier of the QoS policy. Valid in POST and PATCH.

### guarantee

Properties that request and report the space guarantee for the LUN.

Name	Type	Description
requested	boolean	The requested space reservation policy for the LUN. If <i>true</i> , a space reservation is requested for the LUN; if <i>false</i> , 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	Reports if the LUN is space guaranteed.  If <i>true</i> , a space guarantee is requested and the containing volume and aggregate support the request. If <i>false</i> , 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

The storage space related properties of the LUN.

Name	Type	Description
guarantee	guarantee	Properties that request and report the space guarantee for the LUN.
size	integer	<p>The total provisioned size of the LUN. The LUN size can be increased but not be made smaller using the REST interface.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• example: 1073741824</li> </ul>
used	integer	<p>The amount of space consumed by the main data stream of the LUN.</p> <p>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.</p> <p>For more information, see <i>Size properties</i> in the <i>docs</i> section of the ONTAP REST API documentation.</p> <ul style="list-style-type: none"> <li>• readOnly: 1</li> </ul>

status

Status information about the LUN.

Name	Type	Description
container_state	string	The state of the volume and aggregate that contain the LUN. LUNs are only available when their containers are available.
mapped	boolean	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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.
read_only	boolean	Reports if the LUN allows only read access.
state	string	The state of the LUN. Normal states for a LUN are <i>online</i> and <i>offline</i> . Other states indicate errors.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

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.

Name	Type	Description
_links	<a href="#">_links</a>	
auto_delete	boolean	<p>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.</p> <p>When set to <i>true</i>, 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.</p> <p>This property is optional in POST and PATCH. The default value for a new LUN is <i>false</i>.</p> <p>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 <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>
class	string	The class of LUN. Only <i>regular</i> LUNs can be created using the REST API.

Name	Type	Description
clone	clone	<p>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: <code>auto_delete</code>, <code>qos_policy</code>, and <code>space.guarantee.requested</code>.</p> <p>When used in a PATCH, the patched LUN's data is overwritten as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: <code>class</code>, <code>auto_delete</code>, <code>lun_maps</code>, <code>serial_number</code>, <code>status.state</code>, and <code>uuid</code>.</p> <p>Persistent reservations for the patched LUN are also preserved.</p>
comment	string	A configurable comment available for use by the administrator. Valid in POST and PATCH.
enabled	boolean	<p>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 <code>state</code> property to determine if the LUN is administratively disabled (<i>offline</i>) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the <code>enabled</code> property to <i>true</i> or brought administratively offline by setting the <code>enabled</code> property to <i>false</i>. Upon creation, a LUN is enabled by default. Valid in PATCH.</p>

Name	Type	Description
location	<a href="#">location</a>	The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.
lun_maps	array[ <a href="#">lun_maps</a> ]	<p>The LUN maps with which the LUN is associated.</p> <p>There is an added cost to retrieving property values for <code>lun_maps</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>

Name	Type	Description
movement	movement	<p>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 <code>name</code>, <code>location.volume.uuid</code>, or <code>location.volume.name</code>. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.</p> <p>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 <code>movement</code> sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.</p> <p>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 <code>movement</code> properties. The LUN movement operation can be further modified using a PATCH on the properties on the <code>movement</code> sub-object.</p> <p>There is added cost to retrieving property values for <code>movement</code>. They are not populated for either a collection GET or an instance GET unless explicitly requested using the <code>fields</code> query parameter. See <a href="#">DOC Requesting specific fields</a> to learn more.</p>



Name	Type	Description
name	string	<p>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.</p> <p>A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.</p> <p>A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation.</p>
os_type	string	<p>The operating system type of the LUN.</p> <p>Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone.</p>
qos_policy	qos_policy	<p>The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property <code>qos_policy.uuid</code> and <code>qos_policy.name</code> 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 <code>qos_policy.name</code> to an empty string ("") in a PATCH request. Valid in POST and PATCH.</p> <p>Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.</p>

Name	Type	Description
serial_number	string	The LUN serial number. The serial number is generated by ONTAP when the LUN is created. <ul style="list-style-type: none"> <li>• maxLength: 12</li> <li>• minLength: 12</li> <li>• readOnly: 1</li> </ul>
space	<a href="#">space</a>	The storage space related properties of the LUN.
status	<a href="#">status</a>	Status information about the LUN.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created. <ul style="list-style-type: none"> <li>• example: 1cd8a442-86d1-11e0-ae1c-123478563412</li> <li>• readOnly: 1</li> </ul>

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## SVM

### 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

### 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.uuid}/{svm.uuid}
- PATCH /api/svm/peer-permissions/{cluster.uuid}/{svm.uuid}
- DELETE /api/svm/peer-permissions/{cluster.uuid}/{svm.uuid}

## Retrieve SVM peer permissions

GET /svm/peer-permissions

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.

1. Retrieves a list of SVM peer permissions of a specific local SVM

```
GET "/api/svm/peer-permissions/?svm.name=VS1"
```

1. Retrieves a list of SVM peer permissions of a specific cluster peer

```
GET "/api/svm/peer-permissions/?cluster_peer.name=cluster2"
```

### Learn more

- [DOC /svm/peer-permissions](#)

### Parameters

Name	Type	In	Required	Description
applications	string	query	False	Filter by applications
cluster_peer.name	string	query	False	Filter by cluster_peer.name
cluster_peer.uuid	string	query	False	Filter by cluster_peer.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.

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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">svm_peer_permission</a> ]	

## Example response

```
{
  "_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.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

cluster\_peer

Peer cluster details

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "\*".

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

svm\_peer\_permission

Manage SVM peer permissions.

Name	Type	Description
_links	<a href="#">_links</a>	



Name	Type	Description
applications	array[string]	A list of applications for an SVM peer relation.
cluster_peer	<a href="#">cluster_peer</a>	Peer cluster details
svm	<a href="#">svm</a>	Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "*".

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Create an SVM peer permission

POST /svm/peer-permissions

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.

1. 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"]}'
```

1. 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"]}'
```

1. 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

- [DOC /svm/peer-permissions](#)

## Request Body

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>applications</code>	<code>array[string]</code>	A list of applications for an SVM peer relation.
<code>cluster_peer</code>	<code>cluster_peer</code>	Peer cluster details
<code>svm</code>	<code>svm</code>	Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as <code>"*"</code> .

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "flexcache"
  ],
  "cluster_peer": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebef2"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

## Response

Status: 201, Created

Name	Type	Description
<code>_links</code>	<code>_links</code>	
applications	array[string]	A list of applications for an SVM peer relation.
cluster_peer	<code>cluster_peer</code>	Peer cluster details

Name	Type	Description
svm	svm	Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "*".

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "flexcache"
  ],
  "cluster_peer": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebef2"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

### Error

Status: Default

### 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.

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	<a href="#">href</a>	

cluster\_peer

Peer cluster details

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "\*" .

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

svm\_peer\_permission

Manage SVM peer permissions.

Name	Type	Description
_links	<a href="#">_links</a>	
applications	array[string]	A list of applications for an SVM peer relation.
cluster_peer	<a href="#">cluster_peer</a>	Peer cluster details

Name	Type	Description
svm	svm	Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "".

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	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 SVM peer permissions

```
DELETE /svm/peer-permissions/{cluster.uuid}/{svm.uuid}
```

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

- [DOC /svm/peer-permissions](#)

## Parameters

Name	Type	In	Required	Description
cluster.uuid	string	path	True	
svm.uuid	string	path	True	

## 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.

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an SVM peer permission instance

```
GET /svm/peer-permissions/{cluster.uuid}/{svm.uuid}
```

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

- [DOC /svm/peer-permissions](#)

## Parameters

Name	Type	In	Required	Description
cluster.uuid	string	path	True	Peer cluster UUID
svm.uuid	string	path	True	SVM UUID
fields	array[string]	query	False	Specify the fields to return.
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.

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
applications	array[string]	A list of applications for an SVM peer relation.
cluster_peer	<a href="#">cluster_peer</a>	Peer cluster details
svm	<a href="#">svm</a>	Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "*".

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "flexcache"
  ],
  "cluster_peer": {
    "_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"
  }
}
```

## Error

Status: Default

## ONTAP Error Response Codes

Error codes	Description
26345574	Failed to find the SVM or volume name with 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
self	<a href="#">href</a>	

cluster\_peer

Peer cluster details

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "\*" .

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments

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

### Update SVM peer permissions

PATCH /svm/peer-permissions/{cluster.uuid}/{svm.uuid}

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

- [DOC /svm/peer-permissions](#)

#### Parameters

Name	Type	In	Required	Description
cluster.uuid	string	path	True	Peer cluster UUID
svm.uuid	string	path	True	SVM UUID

#### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
applications	array[string]	A list of applications for an SVM peer relation.

Name	Type	Description
cluster_peer	<a href="#">cluster_peer</a>	Peer cluster details
svm	<a href="#">svm</a>	Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "*".

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "flexcache"
  ],
  "cluster_peer": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebef2"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
applications	array[string]	A list of applications for an SVM peer relation.
cluster_peer	<a href="#">cluster_peer</a>	Peer cluster details
svm	<a href="#">svm</a>	Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "**".



## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "flexcache"
  ],
  "cluster_peer": {
    "_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"
  }
}
```

## 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.

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	<a href="#">href</a>	

cluster\_peer

Peer cluster details

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "\*" .

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

svm\_peer\_permission

Manage SVM peer permissions.

Name	Type	Description
_links	<a href="#">_links</a>	
applications	array[string]	A list of applications for an SVM peer relation.
cluster_peer	<a href="#">cluster_peer</a>	Peer cluster details

Name	Type	Description
svm	<a href="#">svm</a>	Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "".

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage SVM peer relationships

### 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/{peer.uuid}
- PATCH /api/svm/peers/{peer.uuid}
- DELETE /api/svm/peers/{peer.uuid}

## Retrieve SVM peer relationships

GET /svm/peers

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.

1. Retrieves a list of SVM peers of a specific local SVM

```
GET "/api/svm/peers/?svm.name=VS1"
```

1. Retrieves a list of SVM peers of a specific cluster peer

```
GET "/api/svm/peers/?peer.cluster.name=cluster2"
```

### Learn more

- [DOC /svm/peers](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	query	False	Filter by uuid
applications	string	query	False	Filter by applications
name	string	query	False	Filter by name
state	string	query	False	Filter by state
peer.cluster.name	string	query	False	Filter by peer.cluster.name
peer.cluster.uuid	string	query	False	Filter by peer.cluster.uuid
peer.svm.uuid	string	query	False	Filter by peer.svm.uuid

Name	Type	In	Required	Description
peer.svm.name	string	query	False	Filter by peer.svm.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	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records

Name	Type	Description
records	array[svm_peer]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "applications": [
      "snapmirror",
      "lun_copy"
    ],
    "peer": {
      "cluster": {
        "_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"
      }
    },
    "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.

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

cluster

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

peer

Details for a peer SVM object.

Name	Type	Description
cluster	<a href="#">cluster</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

svm

#### Local SVM details

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

svm\_peer

An SVM peer relation object.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
applications	array[string]	A list of applications for an SVM peer relation.
name	string	A peer SVM alias name to avoid a name conflict on the local cluster.
peer	<a href="#">peer</a>	Details for a peer SVM object.
state	string	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	<a href="#">svm</a>	Local SVM details
uuid	string	SVM peer relation UUID

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create a new SVM peer relationship

POST /svm/peers

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

- [DOC /svm/peers](#)

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
applications	array[string]	A list of applications for an SVM peer relation.
name	string	A peer SVM alias name to avoid a name conflict on the local cluster.
peer	<a href="#">peer</a>	Details for a peer SVM object.
state	string	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	<a href="#">svm</a>	Local SVM details
uuid	string	SVM peer relation UUID

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "lun_copy"
  ],
  "peer": {
    "cluster": {
      "_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"
    }
  },
  "state": "peered",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}
```

## Response

Status: 202, Accepted

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
applications	array[string]	A list of applications for an SVM peer relation.
name	string	A peer SVM alias name to avoid a name conflict on the local cluster.
peer	<a href="#">peer</a>	Details for a peer SVM object.
state	string	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	<a href="#">svm</a>	Local SVM details
uuid	string	SVM peer relation UUID

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "lun_copy"
  ],
  "peer": {
    "cluster": {
      "_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"
    }
  },
  "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
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.

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	<a href="#">href</a>	

cluster

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

peer

Details for a peer SVM object.

Name	Type	Description
cluster	<a href="#">cluster</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

svm

Local SVM details

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

svm\_peer

An SVM peer relation object.

Name	Type	Description
_links	<a href="#">_links</a>	
applications	array[string]	A list of applications for an SVM peer relation.
name	string	A peer SVM alias name to avoid a name conflict on the local cluster.
peer	<a href="#">peer</a>	Details for a peer SVM object.
state	string	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	<a href="#">svm</a>	Local SVM details
uuid	string	SVM peer relation UUID

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete an SVM peer relationship

```
DELETE /svm/peers/{peer.uuid}
```

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

- [DOC /svm/peers](#)

#### Parameters

Name	Type	In	Required	Description
peer.uuid	string	path	True	

#### 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.

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve an SVM peer relationship instance

GET /svm/peers/{peer.uuid}

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

- [DOC /svm/peers](#)

### Parameters

Name	Type	In	Required	Description
peer.uuid	string	path	True	SVM peer relation UUID
fields	array[string]	query	False	Specify the fields to return.
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.

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
applications	array[string]	A list of applications for an SVM peer relation.
name	string	A peer SVM alias name to avoid a name conflict on the local cluster.
peer	<a href="#">peer</a>	Details for a peer SVM object.

Name	Type	Description
state	string	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	<a href="#">svm</a>	Local SVM details
uuid	string	SVM peer relation UUID



## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "lun_copy"
  ],
  "peer": {
    "cluster": {
      "_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"
    }
  },
  "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.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

cluster

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

peer

Details for a peer SVM object.

Name	Type	Description
cluster	<a href="#">cluster</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

svm

Local SVM details

Name	Type	Description
_links	<a href="#">_links</a>	

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update an SVM peer relationship

PATCH /svm/peers/{peer.uuid}

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.

1. Accepts an SVM peer relationship

```
PATCH "/api/svm/peers/d3268a74-ee76-11e8-a9bb-005056ac6dc9"
'{"state":"peered"}'
```

## 1. Updates the local name of an SVM peer relationship

```
PATCH "/api/svm/peers/d3268a74-ee76-11e8-a9bb-005056ac6dc9"  
' {"name": "vs2"} '
```

### Learn more

- [DOC /svm/peers](#)

### Parameters

Name	Type	In	Required	Description
peer.uuid	string	path	True	SVM peer relationship UUID

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
applications	array[string]	A list of applications for an SVM peer relation.
name	string	A peer SVM alias name to avoid a name conflict on the local cluster.
peer	<a href="#">peer</a>	Details for a peer SVM object.
state	string	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	<a href="#">svm</a>	Local SVM details
uuid	string	SVM peer relation UUID

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "lun_copy"
  ],
  "peer": {
    "cluster": {
      "_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"
    }
  },
  "state": "peered",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "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.

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	<a href="#">href</a>	

cluster

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

peer

Details for a peer SVM object.

Name	Type	Description
cluster	<a href="#">cluster</a>	
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

svm

Local SVM details

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

svm\_peer

An SVM peer relation object.

Name	Type	Description
_links	<a href="#">_links</a>	
applications	array[string]	A list of applications for an SVM peer relation.
name	string	A peer SVM alias name to avoid a name conflict on the local cluster.
peer	<a href="#">peer</a>	Details for a peer SVM object.
state	string	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	<a href="#">svm</a>	Local SVM details
uuid	string	SVM peer relation UUID

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage SVMs

### 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:

1. Data access protocol configuration  
Configures data access protocols, such as NFS, CIFS, iSCSI, and Fibre Channel (FC) protocol (Fibre Channel over Ethernet included).
2. Services configuration  
Configures services such as LDAP, NIS, and DNS.
3. Monitoring SVM  
Monitors jobs, network connections, network interfaces, and SVM health.

#### Retrieve SVMs and SVM properties

GET /svm/svms

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.

### Related ONTAP commands

- `vserver show`

### Examples

1. Retrieves a list of SVMs in the cluster sorted by name

```
GET "/api/svm/svms?order_by=name"
```

1. Retrieves a list of SVMs in the cluster that have the NFS protocol enabled

```
GET "/api/svm/svms?nfs.enabled=true"
```

1. Retrieves a list of SVMs in the cluster that have the CIFS protocol enabled

```
GET "/api/svm/svms?cifs.enabled=true"
```

### Learn more

- [DOC /svm/svms](#)

### Parameters

Name	Type	In	Required	Description
nis.enabled	boolean	query	False	Filter by nis.enabled
nis.servers	string	query	False	Filter by nis.servers
nis.domain	string	query	False	Filter by nis.domain
nvme.enabled	boolean	query	False	Filter by nvme.enabled
language	string	query	False	Filter by language
nfs.enabled	boolean	query	False	Filter by nfs.enabled
comment	string	query	False	Filter by comment
aggregates.name	string	query	False	Filter by aggregates.name

<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
aggregates.uuid	string	query	False	Filter by aggregates.uuid
subtype	string	query	False	Filter by subtype
dns.servers	string	query	False	Filter by dns.servers
dns.domains	string	query	False	Filter by dns.domains
fcp.enabled	boolean	query	False	Filter by fcp.enabled
iscsi.enabled	boolean	query	False	Filter by iscsi.enabled
name	string	query	False	Filter by name
ipspace.uuid	string	query	False	Filter by ipspace.uuid
ipspace.name	string	query	False	Filter by ipspace.name
ldap.base_dn	string	query	False	Filter by ldap.base_dn
ldap.servers	string	query	False	Filter by ldap.servers
ldap.enabled	boolean	query	False	Filter by ldap.enabled
ldap.bind_dn	string	query	False	Filter by ldap.bind_dn
ldap.ad_domain	string	query	False	Filter by ldap.ad_domain
uuid	string	query	False	Filter by uuid
cifs.name	string	query	False	Filter by cifs.name
cifs.ad_domain.fqdn	string	query	False	Filter by cifs.ad_domain.fqdn

Name	Type	In	Required	Description
cifs.ad_domain.organizational_unit	string	query	False	Filter by cifs.ad_domain.organizational_unit
cifs.enabled	boolean	query	False	Filter by cifs.enabled
nsswitch.netgroup	string	query	False	Filter by nsswitch.netgroup
nsswitch.group	string	query	False	Filter by nsswitch.group
nsswitch.hosts	string	query	False	Filter by nsswitch.hosts
nsswitch.namemap	string	query	False	Filter by nsswitch.namemap
nsswitch.passwd	string	query	False	Filter by nsswitch.passwd
ip_interfaces.ip.address	string	query	False	Filter by ip_interfaces.ip.address
ip_interfaces.name	string	query	False	Filter by ip_interfaces.name
ip_interfaces.uuid	string	query	False	Filter by ip_interfaces.uuid
snapshot_policy.uuid	string	query	False	Filter by snapshot_policy.uuid
snapshot_policy.name	string	query	False	Filter by snapshot_policy.name
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">svm</a> ]	

## 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"
    },
    "cifs": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ad_domain": {
        "fqdn": "example.com"
      },
      "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"
  ]
},
"fcp": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"ip_interfaces": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"ip": {
  "address": "10.10.10.7",
  "netmask": "24"
},
"location": {
  "broadcast_domain": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "bd1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"home_node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"name": "node1",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"name": "lif1",
"service_policy": "default-management",

```

```
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "ipospace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "iscsi": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "language": "c.utf_8",
  "ldap": {
    "servers": {
    }
  },
  "name": "svm1",
  "nfs": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "nis": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "servers": {
    }
  },
  "nsswitch": {
    "group": {
    },
    "hosts": {
    },
    "namemap": {

```

```

    },
    "netgroup": {
    },
    "passwd": {
    }
  },
  "nvme": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "routes": {
    "destination": {
      "address": "10.10.10.7",
      "family": "ipv4",
      "netmask": "24"
    },
    "gateway": "10.1.1.1"
  },
  "snapshot_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "state": "running",
  "subtype": "default",
  "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

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.
organizational_unit	string	Specifies the organizational unit within the Active Directory domain to associate with the CIFS server.

Name	Type	Description
password	string	The account password used to add this CIFS server to the Active Directory. This is not audited. Valid in POST only.
user	string	The user account used to add this CIFS server to the Active Directory. Valid in POST only.

#### cifs

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
<a href="#">ad_domain</a>	<a href="#">ad_domain</a>	
enabled	boolean	Specifies whether or not the CIFS service is administratively enabled.
name	string	The NetBIOS name of the CIFS server.

#### dns

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>
servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.

#### fcp

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Enable Fiber Channel Protocol (FCP)? Setting to true creates a service if not already created.

#### ip

##### IP information

Name	Type	Description
address	string	IPv4 or IPv6 address
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0).

broadcast\_domain

Broadcast domain UUID along with a readable name.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Name of the broadcast domain, scoped to its IPspace
uuid	string	Broadcast domain UUID

home\_node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

location

Home\_node is optional.

Name	Type	Description
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name.
home_node	<a href="#">home_node</a>	

ip\_interface\_svm

Interface parameters. Name and home\_node are optional.

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
location	<a href="#">location</a>	Home_node is optional.



Name	Type	Description
name	string	The name of the interface (optional).
service_policy	string	Built-in service policies for SVMs.
uuid	string	The UUID that uniquely identifies the interface.

### ipspace

Either the UUID or name may be supplied on input.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

### iscsi

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Enable iSCSI? Setting to true creates a service if not already created.

### ldap

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.
base_dn	string	Specifies the default base DN for all searches.

Name	Type	Description
bind_dn	string	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	Enable LDAP? Setting to true creates a configuration if not already created.
servers	array[string]	

#### nfs

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	Enable NFS? Setting to true creates a service if not already created.

#### nis

Name	Type	Description
_links	<a href="#">_links</a>	
domain	string	The NIS domain to which this configuration belongs.
enabled	boolean	Enable NIS? Setting to true creates a configuration if not already created.
servers	array[string]	A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration.

#### nsswitch

##### Name Service Switch Configuration

Name	Type	Description
group	array[string]	Group sources
hosts	array[string]	Host sources

Name	Type	Description
namemap	array[string]	NameMap sources
netgroup	array[string]	NetGroup sources
passwd	array[string]	Password sources

#### nvme

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Enable NVMe? Setting to true creates a service if not already created.

#### ip\_info

##### IP information

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.

#### network\_route\_for\_svm

Name	Type	Description
destination	<a href="#">ip_info</a>	IP information
gateway	string	The IP address of the gateway router leading to the destination.

#### snapshot\_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates.
cifs	<a href="#">cifs</a>	
comment	string	Comment
dns	<a href="#">dns</a>	
fc	<a href="#">fc</a>	
ip_interfaces	array[ <a href="#">ip_interface_svm</a> ]	IP interfaces for the SVM
ipspace	<a href="#">ipspace</a>	Either the UUID or name may be supplied on input.
iscsi	<a href="#">iscsi</a>	
language	string	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	<a href="#">ldap</a>	
name	string	The name of the SVM.
nfs	<a href="#">nfs</a>	
nis	<a href="#">nis</a>	
nsswitch	<a href="#">nsswitch</a>	Name Service Switch Configuration
nvme	<a href="#">nvme</a>	

Name	Type	Description
routes	array[ <a href="#">network_route_for_svm</a> ]	Optional array of routes for the SVM
snapshot_policy	<a href="#">snapshot_policy</a>	This is a reference to the Snapshot copy policy.
state	string	SVM State
subtype	string	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	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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create and provision an SVM

POST /svm/svms

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 should not be greater than 5. When the sixth SVM POST request is issued, the error message "Maximum allowed SVM jobs exceeded. Wait for the existing SVM jobs to complete and try again." will be returned.

### 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

### 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`

#### Examples

1. Creates an SVM with default "snapshot\_policy"

```
POST "/api/svm/svms" '{"name":"testVs",
"snapshot_policy":{"name":"default"}}'
```

1. Creates an SVM and configures NFS, ISCSI and FCP

```
POST "/api/svm/svms" '{"name":"testVs", "nfs":{"enabled":"true"},
"fcp":{"enabled":"true"}, "iscsi":{"enabled":"true"}}'
```

1. Creates an SVM and configures NVMe

```
POST "/api/svm/svms" '{"name":"testVs", "nvme":{"enabled":"true"}}'
```

1. 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"}}'
```

#### 1. 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"]}]}'
```

#### 1. 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"]}]}'
```

#### 1. 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"]}]}'
```

#### 1. 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"]}]}'
```

#### 1. 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"]}]'
```

**Learn more**

- [DOC /svm/svms](#)

**Parameters**

Name	Type	In	Required	Description
return_timeout	integer	query	False	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.

**Request Body**

Name	Type	Description
_links	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates.
cifs	<a href="#">cifs</a>	
comment	string	Comment
dns	<a href="#">dns</a>	
fcp	<a href="#">fcp</a>	
ip_interfaces	array[ <a href="#">ip_interface_svm</a> ]	IP interfaces for the SVM
ipspace	<a href="#">ipspace</a>	Either the UUID or name may be supplied on input.
iscsi	<a href="#">iscsi</a>	
language	string	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	<a href="#">ldap</a>	
name	string	The name of the SVM.
nfs	<a href="#">nfs</a>	
nis	<a href="#">nis</a>	
nsswitch	<a href="#">nsswitch</a>	Name Service Switch Configuration
nvme	<a href="#">nvme</a>	
routes	array[ <a href="#">network_route_for_svm</a> ]	Optional array of routes for the SVM
snapshot_policy	<a href="#">snapshot_policy</a>	This is a reference to the Snapshot copy policy.
state	string	SVM State

<b>Name</b>	<b>Type</b>	<b>Description</b>
subtype	string	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	The unique identifier of the SVM.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cifs": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ad_domain": {
      "fqdn": "example.com"
    },
    "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"
    ]
  },
  "fcp": {
    "_links": {
```

```

    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7",
      "netmask": "24"
    },
    "location": {
      "broadcast_domain": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "bd1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "home_node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "name": "lif1",
    "service_policy": "default-management",
    "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": {
    "servers": {
    }
  },
  "name": "svml",
  "nfs": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "nis": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "servers": {
    }
  },
  "nsswitch": {
    "group": {
    },
    "hosts": {
    },
    "namemap": {
    },
    "netgroup": {
    },
    "passwd": {
    }
  },
  "nvme": {
    "_links": {
      "self": {
```

```
        "href": "/api/resourcelink"
      }
    },
  },
  "routes": {
    "destination": {
      "address": "10.10.10.7",
      "family": "ipv4",
      "netmask": "24"
    },
    "gateway": "10.1.1.1"
  },
  "snapshot_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "state": "running",
  "subtype": "default",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

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

## Error

Status: Default

## ONTAP Error Response Codes

Error codes	Description
2621580	Vserver name, comment and IPspace are allowed for a destination Vserver.
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 and retry.

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	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

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.
organizational_unit	string	Specifies the organizational unit within the Active Directory domain to associate with the CIFS server.
password	string	The account password used to add this CIFS server to the Active Directory. This is not audited. Valid in POST only.
user	string	The user account used to add this CIFS server to the Active Directory. Valid in POST only.

cifs

Name	Type	Description
_links	<a href="#">_links</a>	
ad_domain	<a href="#">ad_domain</a>	
enabled	boolean	Specifies whether or not the CIFS service is administratively enabled.
name	string	The NetBIOS name of the CIFS server.

## dns

Name	Type	Description
_links	<a href="#">_links</a>	
domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>

Name	Type	Description
servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.

#### fcp

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Enable Fiber Channel Protocol (FCP)? Setting to true creates a service if not already created.

#### ip

##### IP information

Name	Type	Description
address	string	IPv4 or IPv6 address
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0).

#### broadcast\_domain

Broadcast domain UUID along with a readable name.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	Name of the broadcast domain, scoped to its IPspace
uuid	string	Broadcast domain UUID

#### home\_node

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

#### location

Home\_node is optional.

Name	Type	Description
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name.
home_node	<a href="#">home_node</a>	

ip\_interface\_svm

Interface parameters. Name and home\_node are optional.

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
location	<a href="#">location</a>	Home_node is optional.
name	string	The name of the interface (optional).
service_policy	string	Built-in service policies for SVMs.
uuid	string	The UUID that uniquely identifies the interface.

ipspace

Either the UUID or name may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

iscsi

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	Enable iSCSI? Setting to true creates a service if not already created.

## ldap

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.
base_dn	string	Specifies the default base DN for all searches.
bind_dn	string	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	Enable LDAP? Setting to true creates a configuration if not already created.
servers	array[string]	

## nfs

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Enable NFS? Setting to true creates a service if not already created.

## nis

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
domain	string	The NIS domain to which this configuration belongs.
enabled	boolean	Enable NIS? Setting to true creates a configuration if not already created.

Name	Type	Description
servers	array[string]	A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration.

nsswitch

Name Service Switch Configuration

Name	Type	Description
group	array[string]	Group sources
hosts	array[string]	Host sources
namemap	array[string]	NameMap sources
netgroup	array[string]	NetGroup sources
passwd	array[string]	Password sources

nvme

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Enable NVMe? Setting to true creates a service if not already created.

ip\_info

IP information

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.

network\_route\_for\_svm

Name	Type	Description
destination	<a href="#">ip_info</a>	IP information
gateway	string	The IP address of the gateway router leading to the destination.

snapshot\_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates.
cifs	<a href="#">cifs</a>	
comment	string	Comment
dns	<a href="#">dns</a>	
fcg	<a href="#">fcg</a>	
ip_interfaces	array[ <a href="#">ip_interface_svm</a> ]	IP interfaces for the SVM
ipspace	<a href="#">ipspace</a>	Either the UUID or name may be supplied on input.
iscsi	<a href="#">iscsi</a>	
language	string	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	<a href="#">ldap</a>	



Name	Type	Description
name	string	The name of the SVM.
nfs	<a href="#">nfs</a>	
nis	<a href="#">nis</a>	
nsswitch	<a href="#">nsswitch</a>	Name Service Switch Configuration
nvme	<a href="#">nvme</a>	
routes	array[ <a href="#">network_route_for_svm</a> ]	Optional array of routes for the SVM
snapshot_policy	<a href="#">snapshot_policy</a>	This is a reference to the Snapshot copy policy.
state	string	SVM State
subtype	string	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	The unique identifier of the SVM.

#### job\_link

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete an SVM

```
DELETE /svm/svms/{uuid}
```

Deletes an SVM. As a prerequisite, SVM objects must be deleted first. SnapMirror relations must be deleted and data volumes must be offline and deleted.

The number of parallel SVMs that can be deleted should not be greater than 5. When the sixth SVM DELETE request is issued, the error message "Maximum allowed SVM jobs exceeded. Wait for the existing SVM jobs to complete and try again." will be returned.

### Related ONTAP commands

- `vserver delete`

### Example

Deletes an individual SVM in the cluster.

```
DELETE "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
```

### Learn more

- [DOC /svm/svms](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Filter by UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	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.

### Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

```
{
  "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.

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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve SVM properties

GET /svm/svms/{uuid}

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 note

- 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.

### Example

Retrieves an individual SVM in the cluster

```
GET "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
```

### Learn more

- [DOC /svm/svms](#)

### Parameters

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

### Response

```
Status: 200, Ok
```

Name	Type	Description
_links	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates.
cifs	<a href="#">cifs</a>	
comment	string	Comment
dns	<a href="#">dns</a>	

Name	Type	Description
fcg	fcg	
ip_interfaces	array[ip_interface_svm]	IP interfaces for the SVM
ipspace	ipspace	Either the UUID or name may be supplied on input.
iscsi	iscsi	
language	string	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	ldap	
name	string	The name of the SVM.
nfs	nfs	
nis	nis	
nsswitch	nsswitch	Name Service Switch Configuration
nvme	nvme	
routes	array[network_route_for_svm]	Optional array of routes for the SVM
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.
state	string	SVM State
subtype	string	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	The unique identifier of the SVM.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cifs": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ad_domain": {
      "fqdn": "example.com"
    },
    "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"
    ]
  },
  "fcp": {
    "_links": {
```



```

    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7",
      "netmask": "24"
    },
    "location": {
      "broadcast_domain": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "bd1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "home_node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "name": "lif1",
    "service_policy": "default-management",
    "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": {
  "servers": {
  }
},
"name": "svml",
"nfs": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"nis": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "servers": {
  }
},
"nsswitch": {
  "group": {
  },
  "hosts": {
  },
  "namemap": {
  },
  "netgroup": {
  },
  "passwd": {
  }
},
"nvme": {
  "_links": {
    "self": {
```

```
    "href": "/api/resourcelink"
  }
}
},
"routes": {
  "destination": {
    "address": "10.10.10.7",
    "family": "ipv4",
    "netmask": "24"
  },
  "gateway": "10.1.1.1"
},
"snapshot_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"state": "running",
"subtype": "default",
"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	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

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.
organizational_unit	string	Specifies the organizational unit within the Active Directory domain to associate with the CIFS server.
password	string	The account password used to add this CIFS server to the Active Directory. This is not audited. Valid in POST only.
user	string	The user account used to add this CIFS server to the Active Directory. Valid in POST only.

cifs

Name	Type	Description
_links	<a href="#">_links</a>	
ad_domain	<a href="#">ad_domain</a>	
enabled	boolean	Specifies whether or not the CIFS service is administratively enabled.
name	string	The NetBIOS name of the CIFS server.

#### dns

Name	Type	Description
_links	<a href="#">_links</a>	
domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>

Name	Type	Description
servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.

#### fcv

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Enable Fiber Channel Protocol (FCP)? Setting to true creates a service if not already created.

#### ip

##### IP information

Name	Type	Description
address	string	IPv4 or IPv6 address
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0).

#### broadcast\_domain

Broadcast domain UUID along with a readable name.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	Name of the broadcast domain, scoped to its IPspace
uuid	string	Broadcast domain UUID

#### home\_node

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

#### location

Home\_node is optional.

Name	Type	Description
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name.
home_node	<a href="#">home_node</a>	

ip\_interface\_svm

Interface parameters. Name and home\_node are optional.

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
location	<a href="#">location</a>	Home_node is optional.
name	string	The name of the interface (optional).
service_policy	string	Built-in service policies for SVMs.
uuid	string	The UUID that uniquely identifies the interface.

ipspace

Either the UUID or name may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

iscsi

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	Enable iSCSI? Setting to true creates a service if not already created.



## ldap

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.
base_dn	string	Specifies the default base DN for all searches.
bind_dn	string	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	Enable LDAP? Setting to true creates a configuration if not already created.
servers	array[string]	

## nfs

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Enable NFS? Setting to true creates a service if not already created.

## nis

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
domain	string	The NIS domain to which this configuration belongs.
enabled	boolean	Enable NIS? Setting to true creates a configuration if not already created.

Name	Type	Description
servers	array[string]	A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration.

nsswitch

Name Service Switch Configuration

Name	Type	Description
group	array[string]	Group sources
hosts	array[string]	Host sources
namemap	array[string]	NameMap sources
netgroup	array[string]	NetGroup sources
passwd	array[string]	Password sources

nvme

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Enable NVMe? Setting to true creates a service if not already created.

ip\_info

IP information

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.

network\_route\_for\_svm

Name	Type	Description
destination	<a href="#">ip_info</a>	IP information
gateway	string	The IP address of the gateway router leading to the destination.

snapshot\_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update SVM properties

```
PATCH /svm/svms/{uuid}
```

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`

## Examples

1. 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."}'
```

1. 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."}'
```

1. Updates the "language" field for an individual SVM

```
PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"language":"en.UTF-8"}'
```

1. Updates the "name" field for an SVM or renames the SVM

```
PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"name":"svm_new"}'
```

1. Updates the aggregates for an individual SVM

```
PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"aggregates":{"name":["aggr1", "aggr2", "aggr3"]}}'
```

1. Updates the Snapshot copy policy for an individual SVM

```
PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"snapshot_policy":{"name":"custom1"}}'
```

## Learn more

- [DOC /svm/svms](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Filter by UUID
return_timeout	integer	query	False	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.

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates.
cifs	<a href="#">cifs</a>	
comment	string	Comment

Name	Type	Description
dns	<a href="#">dns</a>	
fcp	<a href="#">fcp</a>	
ip_interfaces	array[ <a href="#">ip_interface_svm</a> ]	IP interfaces for the SVM
ipspace	<a href="#">ipspace</a>	Either the UUID or name may be supplied on input.
iscsi	<a href="#">iscsi</a>	
language	string	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	<a href="#">ldap</a>	
name	string	The name of the SVM.
nfs	<a href="#">nfs</a>	
nis	<a href="#">nis</a>	
nsswitch	<a href="#">nsswitch</a>	Name Service Switch Configuration
nvme	<a href="#">nvme</a>	
routes	array[ <a href="#">network_route_for_svm</a> ]	Optional array of routes for the SVM
snapshot_policy	<a href="#">snapshot_policy</a>	This is a reference to the Snapshot copy policy.
state	string	SVM State
subtype	string	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	The unique identifier of the SVM.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cifs": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ad_domain": {
      "fqdn": "example.com"
    },
    "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"
    ]
  },
  "fcp": {
    "_links": {
```

```

    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7",
      "netmask": "24"
    },
    "location": {
      "broadcast_domain": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "bd1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "home_node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "name": "lif1",
    "service_policy": "default-management",
    "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": {
  "servers": {
  }
},
"name": "svml",
"nfs": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"nis": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "servers": {
  }
},
"nsswitch": {
  "group": {
  },
  "hosts": {
  },
  "namemap": {
  },
  "netgroup": {
  },
  "passwd": {
  }
},
"nvme": {
  "_links": {
    "self": {
```

```
        "href": "/api/resourcelink"
      }
    },
  },
  "routes": {
    "destination": {
      "address": "10.10.10.7",
      "family": "ipv4",
      "netmask": "24"
    },
    "gateway": "10.1.1.1"
  },
  "snapshot_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "state": "running",
  "subtype": "default",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

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

## Error

Status: Default

## 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.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

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.
organizational_unit	string	Specifies the organizational unit within the Active Directory domain to associate with the CIFS server.
password	string	The account password used to add this CIFS server to the Active Directory. This is not audited. Valid in POST only.
user	string	The user account used to add this CIFS server to the Active Directory. Valid in POST only.

cifs

Name	Type	Description
_links	<a href="#">_links</a>	
ad_domain	<a href="#">ad_domain</a>	
enabled	boolean	Specifies whether or not the CIFS service is administratively enabled.
name	string	The NetBIOS name of the CIFS server.

## dns

Name	Type	Description
_links	<a href="#">_links</a>	
domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> <li>• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".</li> <li>• 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.</li> <li>• 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.</li> <li>• The top level domain must contain only the following characters: A through Z, a through z.</li> <li>• The system reserves the following names: "all", "local", and "localhost".</li> </ul>

Name	Type	Description
servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.

#### fcv

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Enable Fiber Channel Protocol (FCP)? Setting to true creates a service if not already created.

#### ip

##### IP information

Name	Type	Description
address	string	IPv4 or IPv6 address
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0).

#### broadcast\_domain

Broadcast domain UUID along with a readable name.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	Name of the broadcast domain, scoped to its IPspace
uuid	string	Broadcast domain UUID

#### home\_node

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

#### location

Home\_node is optional.

Name	Type	Description
broadcast_domain	<a href="#">broadcast_domain</a>	Broadcast domain UUID along with a readable name.
home_node	<a href="#">home_node</a>	

ip\_interface\_svm

Interface parameters. Name and home\_node are optional.

Name	Type	Description
_links	<a href="#">_links</a>	
ip	<a href="#">ip</a>	IP information
location	<a href="#">location</a>	Home_node is optional.
name	string	The name of the interface (optional).
service_policy	string	Built-in service policies for SVMs.
uuid	string	The UUID that uniquely identifies the interface.

ipspace

Either the UUID or name may be supplied on input.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	IPspace name
uuid	string	IPspace UUID

iscsi

Name	Type	Description
_links	<a href="#">_links</a>	
enabled	boolean	Enable iSCSI? Setting to true creates a service if not already created.



## ldap

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.
base_dn	string	Specifies the default base DN for all searches.
bind_dn	string	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	Enable LDAP? Setting to true creates a configuration if not already created.
servers	array[string]	

## nfs

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Enable NFS? Setting to true creates a service if not already created.

## nis

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
domain	string	The NIS domain to which this configuration belongs.
enabled	boolean	Enable NIS? Setting to true creates a configuration if not already created.

Name	Type	Description
servers	array[string]	A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration.

nsswitch

Name Service Switch Configuration

Name	Type	Description
group	array[string]	Group sources
hosts	array[string]	Host sources
namemap	array[string]	NameMap sources
netgroup	array[string]	NetGroup sources
passwd	array[string]	Password sources

nvme

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
enabled	boolean	Enable NVMe? Setting to true creates a service if not already created.

ip\_info

IP information

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.

network\_route\_for\_svm

Name	Type	Description
destination	<a href="#">ip_info</a>	IP information
gateway	string	The IP address of the gateway router leading to the destination.

snapshot\_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

svm

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates.
cifs	<a href="#">cifs</a>	
comment	string	Comment
dns	<a href="#">dns</a>	
fcg	<a href="#">fcg</a>	
ip_interfaces	array[ <a href="#">ip_interface_svm</a> ]	IP interfaces for the SVM
ipspace	<a href="#">ipspace</a>	Either the UUID or name may be supplied on input.
iscsi	<a href="#">iscsi</a>	
language	string	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	<a href="#">ldap</a>	

Name	Type	Description
name	string	The name of the SVM.
nfs	<a href="#">nfs</a>	
nis	<a href="#">nis</a>	
nsswitch	<a href="#">nsswitch</a>	Name Service Switch Configuration
nvme	<a href="#">nvme</a>	
routes	array[ <a href="#">network_route_for_svm</a> ]	Optional array of routes for the SVM
snapshot_policy	<a href="#">snapshot_policy</a>	This is a reference to the Snapshot copy policy.
state	string	SVM State
subtype	string	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	The unique identifier of the SVM.

#### job\_link

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Security

### 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 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, shelf-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, 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, cert
ONTAPI	password, domain, nsswitch, cert
SSH	password, publickey (key pair), domain, nsswitch



In the above table, "cert" 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.uid 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 -k -u <cluster_admin>:<password> -X POST "https://<mgmt-
ip>/api/security/accounts" -d
'{"name":"cluster_user1","applications":[{"application":"ssh","authentica
tion_methods":["password"],"second_authentication_method":"none"}, {"applica
tion":"http","authentication_methods":["password"]}], "role":"admin", "passw
ord":"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 -k -u <cluster_admin>:<password> -X POST "https://<mgmt-
ip>/api/security/accounts" -d '{"owner":{"uuid":"aaef7c38-4bd3-11e9-b238-
0050568e2e25"},"name":"svm_user1","applications":[{"application":"ssh","au
thentication_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 -k -u <cluster_admin>:<password> -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": "svm1",
      "_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 -k -u <cluster_admin>:<password> -X GET "https://<mgmt-
ip>/api/security/accounts/?scope=cluster"

# The scoped call to retrieve the configured SVM-scoped user accounts:
curl -k -u <cluster_admin>:<password> -X GET "https://<mgmt-
ip>/api/security/accounts/?scope=svm"

# The scoped call to retrieve the user accounts configured for the SVM
"svml":
curl -k -u <cluster_admin>:<password> -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 -k -u <cluster_admin>:<password> -X GET "https://<mgmt-
ip>/api/security/accounts/?role=admin"

```

## Retrieve user accounts in the cluster

GET /security/accounts

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
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">account</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "application": "console",
    "authentication_methods": {
    },
    "second_authentication_method": "none"
  },
  "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": "cluster"
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

role\_reference

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
name	string	Role name

#### account

Name	Type	Description
_links	<a href="#">_links</a>	
applications	array[ <a href="#">account_application</a> ]	
comment	string	Optional comment for the user account.
locked	boolean	Locked status of the account.
name	string	User or group account name
owner	<a href="#">owner</a>	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	<a href="#">role_reference</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments

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

## Create a new user account

POST /security/accounts

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](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>applications</code>	array[ <a href="#">account_application</a> ]	



Name	Type	Description
comment	string	Optional comment for the user account.
locked	boolean	Locked status of the account.
name	string	User or group account name
owner	<a href="#">owner</a>	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	<a href="#">role_reference</a>	
scope	string	Scope of the entity. set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": {
    "application": "console",
    "authentication_methods": {
    },
    "second_authentication_method": "none"
  },
  "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": "cluster"
}
```

## Response

Status: 201, Created

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
7077897	Invalid character in username.
7077898	The username must contain both letters and numbers.
7077899	Username does not meet length requirements.
7077906	A role with that name has not been defined for the Vserver.
7077918	Password cannot contain the username.
7077919	Minimum length for new password does not meet the policy.
7077920	New password must have both letters and numbers.
7077921	Minimum number of special characters required do not meet the policy.
7077929	Cannot lock user with non-password authentication method.
7077940	Password exceeds maximum supported length.
7077941	The defined password composition exceeds the maximum password length of 128 characters.
7078900	The admin password is not set. Set the password by including it in the request.
5636099	User creation with non admin role is not supported for service-processor application.
5636121	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 nswitch.
5636156	The same value cannot be specified for the second-authentication-method and the authentication-method.

Error Code	Description
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.

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	<a href="#">href</a>	

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	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

role\_reference

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Role name

account

Name	Type	Description
_links	<a href="#">_links</a>	
applications	array[ <a href="#">account_application</a> ]	
comment	string	Optional comment for the user account.
locked	boolean	Locked status of the account.
name	string	User or group account name
owner	<a href="#">owner</a>	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	<a href="#">role_reference</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
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 GET call 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 -k -u <cluster_admin>:<password> -X GET "https://<mgmt-
ip>/api/security/accounts/ae7c38-4bd3-11e9-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 -k -u <cluster-admin>:<password> -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":
"vsadmin-backup"}'

# The call to update only the role:
curl -k -u <cluster-admin>:<password> -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 -k -u <cluster-admin>:<password> -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 -k -u <cluster-admin>:<password> -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 -k -u <cluster_admin>:<password> -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}

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	
name	string	path	True	

### Response

```
Status: 200, Ok
```

### Error

```
Status: Default
```

### ONTAP Error Response Codes

Error Code	Description
5636098	The last unlocked account that has an admin role cannot be deleted.
5636125	Operation 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[ <a href="#">error_arguments</a> ]	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}

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	
name	string	path	True	

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">account</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "application": "console",
    "authentication_methods": {
    },
    "second_authentication_method": "none"
  },
  "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": "cluster"
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

role\_reference

Name	Type	Description
_links	<a href="#">_links</a>	



Name	Type	Description
name	string	Role name

#### account

Name	Type	Description
_links	<a href="#">_links</a>	
applications	array[ <a href="#">account_application</a> ]	
comment	string	Optional comment for the user account.
locked	boolean	Locked status of the account.
name	string	User or group account name
owner	<a href="#">owner</a>	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	<a href="#">role_reference</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments

Name	Type	Description
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}`

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	Owner UUID of the account.
name	string	path	True	User account name

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
applications	array[ <a href="#">account_application</a> ]	
comment	string	Optional comment for the user account.
locked	boolean	Locked status of the account.
name	string	User or group account name
owner	<a href="#">owner</a>	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	<a href="#">role_reference</a>	
scope	string	Scope of the entity. set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": {
    "application": "console",
    "authentication_methods": {
    },
    "second_authentication_method": "none"
  },
  "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": "cluster"
}
```

## Response

Status: 200, Ok

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
7077906	A role with that name has not been defined for the Vserver.
7077918	Password cannot contain the username.
7077919	Minimum length for new password does not meet the policy.
7077920	New password must have both letters and numbers.
7077921	Minimum number of special characters required do not meet the policy.
7077929	Cannot lock user with non-password authentication method.
7077940	Password exceeds maximum supported length.
7077941	The defined password composition exceeds the maximum password length of 128 characters.
7078900	The admin password is not set. Set the password by including it in the request.
7077911	The user is not configured to use the password authentication method.
7077896	Cannot lock the account of the last console admin user.
7077924	New password must be different than last N passwords.
7077925	New password must be different to the old password.
5636096	Cannot perform the operation for this user account since the password is not set.
5636097	Operation for User account failed since user password is not set.
5636100	User modification is not supported for service-processor application.
5636125	Operation not supported on autosupport user account which is reserved.
5636129	Role does not exist.
5636159	For a given user and application, if the second-authentication-method is specified, only one such login entry is supported.

Error Code	Description
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.
5636174	You are not authorized to change the password for other users.

Name	Type	Description
error	<a href="#">error</a>	

### 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	<a href="#">href</a>	

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	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

role\_reference

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Role name

account

Name	Type	Description
_links	<a href="#">_links</a>	
applications	array[ <a href="#">account_application</a> ]	
comment	string	Optional comment for the user account.
locked	boolean	Locked status of the account.
name	string	User or group account name
owner	<a href="#">owner</a>	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	<a href="#">role_reference</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message



Name	Type	Description
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 operation updates administrative audit settings for GET operations. All fields are optional for the PATCH operation. The GET operation retrieves administrative audit settings for GET operations.

+

#### Examples

##### Retrieving administrative audit settings for GET operations

The following example shows the administrative audit settings for GET operations

+

```
# 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 operations

The following example updates the administrative audit settings for GET operations

+

```
# The API:
/api/security/audit

# The call:
curl -X PATCH "https://<cluster-ip>/api/security/audit" -d
'{"cli":"false", "http": "true", "ontapi": "true"}'
```

## Retrieve the administrative audit settings for GET requests

GET /security/audit

Retrieves administrative audit settings for GET operations.

Learn more

- [DOC /security/audit](#)

### Parameters

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

### 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

## 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the administrative audit settings for GET requests

PATCH `/security/audit`

Updates administrative audit settings for GET operations.  
All of the fields are optional. An empty body will make no changes.

### Learn more

- [DOC /security/audit](#)

### 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	Enable auditing of ONTAP API GET operations. Valid in PATCH

## 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

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## Example error

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

## 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

### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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":
"<destination-address>", "port": <destination-port>, "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/<destination-address>/<destination-
port>"

# 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/<destination-address>/<destination-
port> -d '{"facility": "kern"}'"
```

## 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/<destination-address>/<destination-
port>"
```

## Define a remote syslog or splunk server to receive audit information

GET /security/audit/destinations

Defines remote syslog/splunk server for sending audit information

### Learn more

- [DOC /security/audit/destinations](#)

### Parameters

Name	Type	In	Required	Description
protocol	string	query	False	Filter by protocol
port	integer	query	False	Filter by port
verify_server	boolean	query	False	Filter by verify_server
facility	string	query	False	Filter by facility
address	string	query	False	Filter by address
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	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.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">security_audit_log_forward</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "facility": "kern",
    "protocol": "udp_unencrypted"
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Define the remote syslog or splunk server information

POST `/security/audit/destinations`

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")

#### Learn more

- [DOC /security/audit/destinations](#)

#### Parameters

Name	Type	In	Required	Description
force	boolean	query	False	Skip the Connectivity Test  • 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

```
{
  "facility": "kern",
  "protocol": "udp_unencrypted"
}
```

## Response

Status: 202, Accepted

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">security_audit_log_forward</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "facility": "kern",
    "protocol": "udp_unencrypted"
  }
}
```

## 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.



Error Code	Description
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.

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

### 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

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete the remote syslog or splunk server information

DELETE /security/audit/destinations/{address}/{port}

Deletes remote syslog/splunk server information.

Learn more

- [DOC /security/audit/destinations](#)

### Parameters

Name	Type	In	Required	Description
address	string	path	True	
port	integer	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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Retrieve the remote syslog or splunk server information

GET /security/audit/destinations/{address}/{port}

Defines remote syslog/splunk server for sending audit information.

### Learn more

- [DOC /security/audit/destinations](#)

### 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

```
{
  "facility": "kern",
  "protocol": "udp_unencrypted"
}
```

### 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": "uid"
  }
}
```

## Definitions

### See Definitions

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the remote syslog or splunk server information

PATCH /security/audit/destinations/{address}/{port}

Updates remote syslog/splunk server information.

### Learn more

- [DOC /security/audit/destinations](#)

### 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
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

```
{  
  "facility": "kern",  
  "protocol": "udp_unencrypted"  
}
```

## 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
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": "kern",
  "protocol": "udp_unencrypted"
}
```

### Error

```
Status: Default, Default
```

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

### 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[ <a href="#">error_arguments</a> ]	Message arguments

Name	Type	Description
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 operation retrieves all the 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

Retrieves the administrative audit log viewer.

## Learn more

- [DOC /security/audit/messages](#)

## Parameters

Name	Type	In	Required	Description
user	string	query	False	Filter by user
index	integer	query	False	Filter by index
session_id	string	query	False	Filter by session_id
scope	string	query	False	Filter by scope
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
state	string	query	False	Filter by state
input	string	query	False	Filter by input
location	string	query	False	Filter by location
command_id	string	query	False	Filter by command_id
application	string	query	False	Filter by application
timestamp	string	query	False	Filter by timestamp
message	string	query	False	Filter by 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	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.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">security_audit_log</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "application": "internal",
    "command_id": "string",
    "index": 0,
    "input": "string",
    "location": "string",
    "message": "string",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "scope": "svm",
    "session_id": "string",
    "state": "pending",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

node

Node where the audit message resides.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">node</a>	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	<a href="#">svm</a>	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[ <a href="#">error_arguments</a> ]	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 operation retrieves the LDAP configuration of the cluster.

The following example shows how a GET operation 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\" ], \"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 operation updates the LDAP configuration of the cluster.

The following example shows how a PATCH operation 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 operation deletes the LDAP configuration of the cluster.

The following example shows how a DELETE operation 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

The DELETE operation removes the LDAP configuration of the cluster.

#### Learn more

- [DOC /security/authentication/cluster/ldap](#)

#### Response

```
Status: 200, Ok
```

#### Error

```
Status: Default, Error
```

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	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

Retrieves the cluster LDAP configuration.



## Learn more

- [DOC /security/authentication/cluster/ldap](#)

## Parameters

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

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"><li>• base - search the named entry only</li><li>• onelevel - search all entries immediately below the DN</li><li>• subtree - search the named DN entry and the entire subtree below the DN</li></ul>
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"><li>• anonymous - anonymous bind</li><li>• simple - simple bind</li><li>• sasl - Simple Authentication and Security Layer (SASL) bind</li></ul>

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"> <li>• AD-IDMU - Active Directory Identity Management for UNIX</li> <li>• AD-SFU - Active Directory Services for UNIX</li> <li>• MS-AD-BIS - Active Directory Identity Management for UNIX</li> <li>• RFC-2307 - Schema based on RFC 2307</li> <li>• Custom schema</li> </ul>
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"> <li>• none - no signing or sealing</li> <li>• sign - sign LDAP traffic</li> <li>• seal - seal and sign LDAP traffic</li> </ul>
use_start_tls	boolean	Specifies whether or not to use Start TLS over LDAP connections.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "base_scope": "base",
  "min_bind_level": "anonymous",
  "port": 389,
  "servers": {
  },
  "session_security": "none"
}
```

## 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	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

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' field.

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

## Learn more

- [DOC /security/authentication/cluster/ldap](#)

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• base - search the named entry only</li><li>• onelevel - search all entries immediately below the DN</li><li>• subtree - search the named DN entry and the entire subtree below the DN</li></ul>
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"><li>• anonymous - anonymous bind</li><li>• simple - simple bind</li><li>• sasl - Simple Authentication and Security Layer (SASL) bind</li></ul>
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"> <li>• AD-IDMU - Active Directory Identity Management for UNIX</li> <li>• AD-SFU - Active Directory Services for UNIX</li> <li>• MS-AD-BIS - Active Directory Identity Management for UNIX</li> <li>• RFC-2307 - Schema based on RFC 2307</li> <li>• Custom schema</li> </ul>
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"> <li>• none - no signing or sealing</li> <li>• sign - sign LDAP traffic</li> <li>• seal - seal and sign LDAP traffic</li> </ul>
use_start_tls	boolean	Specifies whether or not to use Start TLS over LDAP connections.

### Example request

```

{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "base_scope": "base",
  "min_bind_level": "anonymous",
  "port": 389,
  "servers": {
  },
  "session_security": "none"
}

```

## 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	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 the Distinguished Names and bind password are correct
4915258	The LDAP configuration is invalid. 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 LIFs

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

cluster\_ldap

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• base - search the named entry only</li><li>• onelevel - search all entries immediately below the DN</li><li>• subtree - search the named DN entry and the entire subtree below the DN</li></ul>
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"><li>• anonymous - anonymous bind</li><li>• simple - simple bind</li><li>• sasl - Simple Authentication and Security Layer (SASL) bind</li></ul>

Name	Type	Description
port	integer	The port used to connect to the LDAP Servers.
schema	string	The name of the schema template used by the SVM. <ul style="list-style-type: none"> <li>• AD-IDMU - Active Directory Identity Management for UNIX</li> <li>• AD-SFU - Active Directory Services for UNIX</li> <li>• MS-AD-BIS - Active Directory Identity Management for UNIX</li> <li>• RFC-2307 - Schema based on RFC 2307</li> <li>• Custom schema</li> </ul>
servers	array[string]	
session_security	string	Specifies the level of security to be used for LDAP communications: <ul style="list-style-type: none"> <li>• none - no signing or sealing</li> <li>• sign - sign LDAP traffic</li> <li>• seal - seal and sign LDAP traffic</li> </ul>
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[ <a href="#">error_arguments</a> ]	Message arguments

Name	Type	Description
code	string	Error code
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`

A cluster can have only one LDAP configuration. IPv6 must be enabled if IPv6 family addresses are specified. The following parameters are optional:

- 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 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.

### Learn more

- [DOC /security/authentication/cluster/ldap](#)

### Request Body

Name	Type	Description
<code>_links</code>	<code><a href="#">_links</a></code>	
<code>base_dn</code>	string	Specifies the default base DN for all searches.

Name	Type	Description
base_scope	string	Specifies the default search scope for LDAP queries: <ul style="list-style-type: none"> <li>• base - search the named entry only</li> <li>• onelevel - search all entries immediately below the DN</li> <li>• subtree - search the named DN entry and the entire subtree below the DN</li> </ul>
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"> <li>• anonymous - anonymous bind</li> <li>• simple - simple bind</li> <li>• sasl - Simple Authentication and Security Layer (SASL) bind</li> </ul>
port	integer	The port used to connect to the LDAP Servers.
schema	string	The name of the schema template used by the SVM. <ul style="list-style-type: none"> <li>• AD-IDMU - Active Directory Identity Management for UNIX</li> <li>• AD-SFU - Active Directory Services for UNIX</li> <li>• MS-AD-BIS - Active Directory Identity Management for UNIX</li> <li>• RFC-2307 - Schema based on RFC 2307</li> <li>• Custom schema</li> </ul>
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"> <li>• none - no signing or sealing</li> <li>• sign - sign LDAP traffic</li> <li>• seal - seal and sign LDAP traffic</li> </ul>
use_start_tls	boolean	Specifies whether or not to use Start TLS over LDAP connections.

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "base_scope": "base",
  "min_bind_level": "anonymous",
  "port": 389,
  "servers": {
  },
  "session_security": "none"
}
```

### Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of LDAP records.
records	array[ <a href="#">ldap_service</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "base_scope": "base",
    "min_bind_level": "anonymous",
    "port": 389,
    "preferred_ad_servers": {
    },
    "servers": {
    },
    "session_security": "none",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "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
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
23724130	Cannot use an IPv6 name server address because there are no IPv6 LIFs

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	<a href="#">href</a>	

cluster\_ldap

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• base - search the named entry only</li><li>• onelevel - search all entries immediately below the DN</li><li>• subtree - search the named DN entry and the entire subtree below the DN</li></ul>
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"><li>• anonymous - anonymous bind</li><li>• simple - simple bind</li><li>• sasl - Simple Authentication and Security Layer (SASL) bind</li></ul>

Name	Type	Description
port	integer	The port used to connect to the LDAP Servers.
schema	string	The name of the schema template used by the SVM. <ul style="list-style-type: none"> <li>• AD-IDMU - Active Directory Identity Management for UNIX</li> <li>• AD-SFU - Active Directory Services for UNIX</li> <li>• MS-AD-BIS - Active Directory Identity Management for UNIX</li> <li>• RFC-2307 - Schema based on RFC 2307</li> <li>• Custom schema</li> </ul>
servers	array[string]	
session_security	string	Specifies the level of security to be used for LDAP communications: <ul style="list-style-type: none"> <li>• none - no signing or sealing</li> <li>• sign - sign LDAP traffic</li> <li>• seal - seal and sign LDAP traffic</li> </ul>
use_start_tls	boolean	Specifies whether or not to use Start TLS over LDAP connections.

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### svm

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.

Name	Type	Description
uuid	string	The unique identifier of the SVM.

#### ldap\_service

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"> <li>• base - search the named entry only</li> <li>• onelevel - search all entries immediately below the DN</li> <li>• subtree - search the named DN entry and the entire subtree below the DN</li> </ul>
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"> <li>• anonymous - anonymous bind</li> <li>• simple - simple bind</li> <li>• sasl - Simple Authentication and Security Layer (SASL) bind</li> </ul>

Name	Type	Description
port	integer	The port used to connect to the LDAP Servers.
preferred_ad_servers	array[string]	
schema	string	<p>The name of the schema template used by the SVM.</p> <ul style="list-style-type: none"> <li>• AD-IDMU - Active Directory Identity Management for UNIX</li> <li>• AD-SFU - Active Directory Services for UNIX</li> <li>• MS-AD-BIS - Active Directory Identity Management for UNIX</li> <li>• RFC-2307 - Schema based on RFC 2307</li> <li>• Custom schema</li> </ul>
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"> <li>• none - no signing or sealing</li> <li>• sign - sign LDAP traffic</li> <li>• seal - seal and sign LDAP traffic</li> </ul>
svm	<a href="#">svm</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
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 operation retrieves the NIS configuration of the cluster.

The following example shows how a GET operation 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 operation creates a NIS configuration for the cluster.

The following example shows how a POST operation 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 operation 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 operation deletes the NIS configuration of the cluster.

The following example shows how a DELETE operation 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

The DELETE operation removes the NIS configuration of the cluster. NIS can be removed as a source from ns-switch if NIS is not used for lookups.

### Learn more

- [DOC /security/authentication/cluster/nis](#)

### Response

Status: 200, Ok

### Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	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

Retrieves the NIS configuration of the cluster. Both NIS domain and servers are displayed by default. The 'bound servers' field indicates the successfully bound NIS servers.



## Learn more

- [DOC /security/authentication/cluster/nis](#)

## Parameters

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

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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": {
  },
  "servers": {
  }
}
```

## 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	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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
```

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.

### Learn more

- [DOC /security/authentication/cluster/nis](#)

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bound_servers": {
  },
  "servers": {
  }
}
```

### Response

Status: 200, Ok

### Error

Status: Default

### ONTAP Error Response Codes

Error Code	Description
1966253	IPv6 is not enabled in the cluster
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

Error Code	Description
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

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	<a href="#">href</a>	

cluster\_nis\_service

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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

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 'servers' field. IPv6 must be enabled if IPv6 family addresses are specified in the 'servers' field.

A maximum of ten NIS servers are supported.

### Learn more

- [DOC /security/authentication/cluster/nis](#)

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bound_servers": {
  },
  "servers": {
  }
}
```

### Response

```
Status: 201, Created
```

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of NIS domain records.
records	array[ <a href="#">cluster_nis_service</a> ]	

### Example response

```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "bound_servers": {
    },
    "servers": {
    }
  }
}

```

### Error

Status: Default

### ONTAP Error Response Codes

Error Code	Description
1966253	IPv6 is not enabled in the cluster
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



Error Code	Description
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

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	<a href="#">href</a>	

cluster\_nis\_service

Name	Type	Description
_links	<a href="#">_links</a>	
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.

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code

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

## Manage SAML service

### Security authentication cluster saml-sp endpoint overview

#### Overview

You can use this API to retrieve and display relevant information pertaining to the SAML service provider configuration in the cluster. The POST operation creates a SAML service provider configuration if there is none present. The DELETE operation removes the SAML service provider configuration. The PATCH operation 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
```

Deletes a SAML service provider configuration.

### Learn more

- [DOC /security/authentication/cluster/saml-sp](#)

### 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	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	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

Retrieves a SAML service provider configuration.

### Learn more

- [DOC /security/authentication/cluster/saml-sp](#)

### Parameters

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

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
certificate	<a href="#">certificate</a>	
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": {
    "common_name": "cluster1",
    "serial_number": "1506B24A94F566BA"
  },
  "idp_uri": "https://idp.example.com/FederationMetadata/2007-06/FederationMetadata.xml"
}
```

### Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	



## 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	<a href="#">href</a>	

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[ <a href="#">error_arguments</a> ]	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

Updates a SAML service provider configuration.

**Learn more**

- [DOC /security/authentication/cluster/saml-sp](#)

**Request Body**

Name	Type	Description
_links	<a href="#">_links</a>	
certificate	<a href="#">certificate</a>	
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 request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "common_name": "cluster1",
    "serial_number": "1506B24A94F566BA"
  },
  "idp_uri": "https://idp.example.com/FederationMetadata/2007-06/FederationMetadata.xml"
}
```

## 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.

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	<a href="#">href</a>	

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
_links	<a href="#">_links</a>	
certificate	<a href="#">certificate</a>	
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.

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

Creates a SAML service provider configuration. Note that "common\_name" is mutually exclusive with "serial\_number" and "ca" in the 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

### Learn more

- [DOC /security/authentication/cluster/saml-sp](#)

### Parameters

Name	Type	In	Required	Description
verify_metadata_server	boolean	query	False	Verify IdP metadata server identity.  • Default value: 1

### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
certificate	<a href="#">certificate</a>	
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 request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "common_name": "cluster1",
    "serial_number": "1506B24A94F566BA"
  },
  "idp_uri": "https://idp.example.com/FederationMetadata/2007-06/FederationMetadata.xml"
}
```

### 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
12320814	An invalid IDP URI has been entered.
12320815	The IDP URI must be an HTTPS or FTPS URI.
12320794	The host parameter provided must be the cluster management LIF's IP address. If the cluster management LIF is not available, the node management LIF's IP address must be used.
12320795	A valid cluster or node management LIF IP address must be provided.
12320805	The certificate information provided does not match any installed certificates.
12320806	Entered certificate information does not match any installed certificates.

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	<a href="#">href</a>	

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
_links	<a href="#">_links</a>	
certificate	<a href="#">certificate</a>	
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.

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update the 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 change your own password only.

#### 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 -k -u <cluster_admin>:<password> -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 -k -u <cluster_admin>:<password> -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 -k -u svm_user1:hello@1234 -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`

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](#)

## Request Body

Name	Type	Description
name	string	The user account name whose password is being modified.
owner	<a href="#">owner</a>	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

```
{
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourceLink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

## Response

Status: 201, Created

## Error

Status: Default

### ONTAP Error Response Codes

Error Code	Description
7077919	Minimum length for new password does not meet the policy.
7077920	New password must have both letters and numbers.
7077921	Minimum number of special characters required do not meet the policy.
7077940	Password exceeds maximum supported length.
7077941	The defined password composition exceeds the maximum password length of 128 characters.
7077918	Password cannot contain the username.
7077924	New password must be different than last N passwords.
7077925	New password must be different to the old password.

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	<a href="#">href</a>	

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
_links	<a href="#">_links</a>	
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	<a href="#">owner</a>	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
message	string	Message argument

## error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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 endpoint 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}
```



## 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-----
\nMIIFYDCCA0igAwIBAgIQcGFCgAAAAUjyESlAAAAAjANBgkqhkiG9w0BAQsFADBKMQswCQYD
VQQG\nEwJVUzESMBAGA1UEChMJSWRlbnRydXN0MScwJQYDVQQDEx5JZGVuVHJlcnQ29tbWVy
Y2lhbCBS\nb290IENBIDEwHhcNMTQwMTE2MTg5MjIzWhcNMzQwMTE2MTg5MjIzWjBKMQswCQYD
VQQGEwJVUzES\nMBAGA1UEChMJSWRlbnRydXN0MScwJQYDVQQDEx5JZGVuVHJlcnQ29tbWVy
Y2lhbCBSb290IENB\nnIDEwggIiMA0GCSqGSIb3DQEBAQUAA4ICDwAwggIKAoICAQCnUBneP5k9
1DNG8W9RYYKyqU+PZ4ld\nhNlT3Qwo2dfw/66VQ3KZ+bVdfIrBQuExUHRgQl8zZshq0PirK1e
hm7zCYofWjK9ouuU+ehcCuz/\nmNKvcb00U59Oh++SvL3sTzIwiEsXXlFEU8L2ApeN2WIrvyQf
Yo3fw7gpS0l4PJNgiCL8mdo2yMKi\nlCcxUAGclbnO/AljwpN3lsKImesrgNqUZFvX9t++uP0D1
bVoE/c40yiTcdCMbXTMTEl3EASX2MN0C\nXZ/g1Ue9tOsobotJSdifWwLziuQkkORiT0/Br4s0
dBeo0XKIanoBScy0RnnGF7HamB4HWfp1IYVl\n3ZBWzvurpWCdxJ35UrClvYf5jysjCiN2O/cz
4ckA82n5S6LgTrx+kzmEB/dEcH7+B1rlsazRGMzy\nNeVJSQjKVsk9+w8YfYs7wRPCTY/JTw43
6R+hDmrfYi7LNQZReSzIJTj0+kuniVyc0uMNOYZkDhZV\nWYfCP04MXFL0PfdSgvHqo6z9STQa
KPNBiDoT7uje/5kdX7rL6B7yuVBgdHTc+XvvqDtMwt0viAg\nxGds8AgDelWAF0ZOlqf0Hj7h
9tgJ4TNkK2PXm16f+cb7D3hvl7yTmvmcEpB4eoCHFddyJxVdHix\nnuuFucAS6T6C6aMN7/zHw
cz09lCqxCOEOoP5NiGVreTO01wIDAQABo0IwQDAOBgNVHQ8BAf8EBAMC\nAQYwDwYDVR0TAQH/
BAUwAwEB/zAdBgNVHQ4EFgQU7UQZwNPwBovupHu+QucmVMiONnYwDQYJKoZI\nnhvcNAQELBQAD
ggIBAA2ukDL2pkt8RHYZYR4nKM1eVO8lvOMIkPkp165oCOGUAFjvLi5+U1KMtlwH\n6oi6mYtQ
lNeCgN9hCQCTrQ0U5s7B8jeUeLBfnLOic7iPBZM4zy0+sLj7wM+x8uwtLRvM7Kqas6pg\nnghst
O8OEPVeKlh6cdbhTMM1gCtIOQ045U8U1mwF10A0Cj7oV+wh93nAbowacYXVKV7cndJZ5t+qnt\n
ozo00F172u1Q8zW/7esUTTHHYPTa8Yec4kjixsU3+wYQ+nVZZjFHKdp2mhZpgq7vmr1R94gJmm
mV\nYjz1VYA211QC//G5Xc7UI2/YRYRKW2XviQzdFKcgyxilJbQN+QHwotL0AMh0jqEqSI5l2x
PE4iUX\nnfeu+h1sXIFRRk0pTAvsXcoz7WL9RccvW9xYoIA55vrX/hMUpu091EpCdNTDd1lzzY
9Gv1U47/ro\nnkTLq11gEIt44w8y8bckzOmoKaT+gyOpyj4xjhi09bTyWnpXgSUyqorkqG5w2gX
jtw+hG4iZZRHUe\n2XWJUc0QhJ1hYmtD+ZciTY6Y5uN/9lu7rs3KSoFrXgvzUeF0K+1+J6fZmU
lO+KWA2yUPHGNiiskz\nZ2s8EIPGrd6ozRaOjfAHN3Gf8qv8QfXBi+wAN10J5U6A7/qxXDgGpR
tK4dw4LTzcx+QGtVKno7R\nncGzM7vRX+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-----
\nMIIFYDCCA0igAwIBAgIQCgFCgAAAAUjyES1AAAAjANBgkqhkiG9w0BAQsFADBKMQswCQYD
VQQG\nEwJVUzESMBAGA1UEChMJSWRlb1RydXN0MScwJQYDVQQDEx5JZGVuVHJlc3QgQ29tbWVy
Y2lhbCBS\nb290IENBIDEwHhcNMTQwMjdfw/66VQ3KZ+bVdfIrBQuExUHTRgQ18zZshq0PirK1e
hm7zCYofWjK9ouuU+ehcCuz/\nmNKvcb00U590h++SvL3sTzIwiEsXXlFEU8L2ApeN2WIrvyQf
Yo3fw7gps014PjNgiCL8mdo2yMKi\n1CxUAGc1bnO/AljwpN3lsKImesrgNqUZFvX9t++uP0D1
bVoE/c40yiTcdCMbXTMTE13EASX2MN0C\nXZ/g1Ue9tOsbobtJSdifWwLziuQkkORiT0/Br4sO
dBeo0XKIanoBScy0RnnGF7HamB4HWfp1IYVl\n3ZBWzvurpWCdxJ35UrCLvYf5jysjCiN20/cz
4ckA82n5S6LgTrx+kzmEB/dEcH7+B1rlsazRGMzy\nNeVJSQjKVsk9+w8YfYs7wRPTY/JTw43
6R+hDmrfYi7LNQZReSzIJTj0+kuniVyc0uMNOYZKdHzV\nWYfCP04MXFL0PfdSgvHqo6z9STQa
KPNBiDoT7uje/5kdX7rL6B7yuVBgwDHTc+XvvqDtMwt0viAg\nxGds8AgDelWaf0ZO1qf0Hj7h
9tgJ4TNkK2PXm16f+cB7D3hvl7yTmvmcEpB4eoCHFddyJxVdHix\nnuuFucAS6T6C6aMN7/zHw
cz09lCqxC0EOoP5NiGVreT001wIDAQBo0IwQDAOBgNVHQ8BAf8EBAMC\nAQYwDwYDVR0TAQH/
BAUwAwEB/zAdBgNVHQ4EFgQU7UQZwNPwBovupHu+QucmVMiONnYwDQYJKoZI\nnhvcNAQELBQAD
ggIBAA2ukDL2pkt8RHYZYR4nKM1eVO81vOMIkPkp165oCOGUAFjvLi5+U1KMtlwH\n6oi6mYtQ
lNeCgN9hCQCTrQ0U5s7B8jeUeLBfnLOic7iPBZM4zY0+sLj7wM+x8uwTLrVm7Kqas6pg\nnghst
O8OEPVeKlh6cdbjTMM1gCIOQ045U8U1mwF10A0Cj7oV+wh93nAbowacYXVKV7cndJZ5t+qnt\n
ozo00F172u1Q8zW/7esUTTHHYPTa8Yec4kjixsU3+wYQ+nVZZjFHKdp2mhzpgq7vmr1R94gjmm
mV\nYjz1VYA211QC//G5Xc7UI2/YRYRKW2XviQzdFKcgyxilJbQN+QHwotL0AMh0jqEqSI512x
PE4iUX\nnfeu+h1sXIFRRk0pTAvvsXcoz7WL9RccvW9xYoIA55vrX/hMUpu091EpCdNTDd1lzzY
9GvlU47/ro\nnkTLq11gEIt44w8y8bckzOmoKaT+gyOpyj4xjhi09bTyWnpXgSUyqorkqG5w2gX
jtw+hG4iZZRHUe\n2XWJUc0QhJ1hYMtD+ZciTY6Y5uN/9lu7rs3KSoFrXgvzUeF0K+l+J6fZmU
lO+KWA2yUPHGNiiskz\n2Zs8EIPGrd6ozRaOjfAHN3Gf8qv8QfXBi+wAN10J5U6A7/qxXDgGpR
tK4dw4LTzcx+QGtVKn07R\nncGzM7vRX+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-----\nMIICYTCCAUKCAQAwHDENMAsGA1UEAxMEVEVTVDELMAkGA1UEBhMCVVMwggEiMA0G\nb3DQEBAQUAA4IBDwAwggEKAoIBAQCIBCuVfbYHNdOO7vjRQja4JqL2cHqK\ndr1Tj5hz9RVqFKZ7VP8DSP9LoTbYWsvrTkbuD0Wi715MVQCsbkq/mHos+Y51fqs\nNP5K92fc6EhBzBDYFgZGFntZYJjEG5MPerIUE7CfVy7o6sjWolxeY33pjefObyvP\nBcJkBHg6SFJK/TDLvIYJkonLkJEOJoTI6++a3I/1bCMfUeuRtLU9ThWlna1kMMYK\n4T16/Bxgm4bha2U2jtosc0Wltnld/capc+eqRV07WVbMmEOTtop3cv0h3N0S61bn\nfkd96DXzeGwbSHFHckeCZ9bOHhnVbfEa/efkPLx7ziMC8GtRHHlwbNk7AgMBAAGg\nADANBgkqhkiG9w0BAQsFAAOCAQEaf+rs1i5PHaOSI2HtTM+Hcv/p71yzgoLL+aeU\nntB0V4iuoXdqY8oQeWoPI92ci0K08JuSpu6D0DwCK1stfwuGkAA2b0Wr7ZDRonT\nUq\nnmJ4j3O47MLysW4Db2LbGws/AuDsCIrBJDWHMpHaqsvRbpMx2xQ/V5oagUw5eGGpN\nne4fg/E2k9mGkpxwUzT7w1RZirpND4xL+XTzpzzeZqgalpXug4yjIXlI5hpRESZ9\n\nAkGJSCWxi15IZdxxFVXlBcmm6WpJnnboqkcKeXz95GM6Re+oBy9tlgvwv1Vd5s8uHX+bycFiZp09Wsm8Ev727MziZ+0II9nxwkDKsdPvam+KLI9hLQ==\n-----END CERTIFICATE REQUEST-----\n\", \"hash_function\": \"sha256\"}"

# The response:
{
  "public_certificate": "-----BEGIN CERTIFICATE-----\nMIIDBzCCAe+gAwIBAgIIFUKQpcqeaUAWDQYJKoZIhvcNAQELBQAwHDENMAsGA1UE\nAxMEUkFDWDELMAkGA1UEBhMCVVMwHhcNMTgwNze4MjAzMTA1WhcNMTkwNze4MjAz\n\nMTA1WjAcMQ0wCwYDVoQDEwRURVNUMQswCQYDVoQGEwJVUzCCASIwDQYJKoZIhvcN\n\nAQEBBQADggEPADCCAQoCggEBAKIEK5V9tgc1047u+NFCNrgmovZweop2uVOPmHP1\n\nFWoUpntU+HwNI/0uhNthay+tORu4PRaLvXkxVAKXuSr+Yeiz5jmV+qw0/kr3Z9zo\n\nnSEHMENgWBkYWellgmMQbkw96shQTsJ9XLujqyNY6XF5jfemN585vK88FwmQEeDpI\n\nnUkr9MMu8hgmSicuQkQ4mhMjr75rcj/VsIx9R65G0tT10FaWdrWQwxgrhPXR8HGCb\n\nnhuFrZTa02ixzRaW2eV39xqlz56pFXTtZVsyYQ502indy/SHc3RLqVucWR33oNfN4\n\nnZZtIcUdyR4Jn1s4eGdVt8Rr95+Q8vHvOIwLwa1EceXBucrsCAwEAAaNNMEswCQYD\n\nnVR0TBAlwADAdBgNVHQ4EFgQUJMPxjeW1G76TbbD2tXB8dwSpI3MwHwYDVR0jBBGw\n\nnFoAUu5aH0mWR4cFoN9i7k96d2op3sPwwDQYJKoZIhvcNAQELBQADggEBAl5ai+Zi\n\nnFQZUXRTqJCgHsgBThARneVWQYkYpyAXmTR7QeLfld4ZHL33i4xWCqX3uvW7SFJLe\n\nnZajT2AVmgiDbaWIHtDtvqz1BY78PSgUwPH/IyARTEOBeikp6KdwMPraehDIBMAcc\n\nnANY58wXiTBbsl8UMD6tGecgnzw6sxlMmadGvrfJeJmgY4zert6NNvgtPPhcZQdLS\n\nnE0fGzHS6+3ajCCfEEhPNPer9D0e5Me81i9EsQGENrnJzTci8rzXPuF4bC3gghrK1\n\nnI1+kmJQ1kLYVUcsntcrIiHmNvtPFJY6stjDgQKS9aDd/THhPpokPtZoCmE6PDxh6\n\nnR+dO6C0hcDKHFzA=\n-----END CERTIFICATE-----\n"
}

```

## Retrieve security certificates

GET /security/certificates

Retrieves security certificates.

### Related ONTAP commands

- `security certificate show`

### Learn more

- [DOC /security/certificates](#)

### Parameters

Name	Type	In	Required	Description
scope	string	query	False	Filter by scope
svm.name	string	query	False	Filter by svm.name
svm.uuid	string	query	False	Filter by svm.uuid
common_name	string	query	False	Filter by common_name
serial_number	string	query	False	Filter by serial_number
ca	string	query	False	Filter by ca
type	string	query	False	Filter by certificate type
key_size	string	query	False	Filter by key_size
expiry_time	string	query	False	Filter by expiry_time
hash_function	string	query	False	Filter by hash_function
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	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.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">security_certificate</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ca": "string",
    "common_name": "test.domain.com",
    "hash_function": "sha1",
    "intermediate_certificates": {
    },
    "private_key": "-----BEGIN PRIVATE KEY-----
MIIBVAIBADANBgkqhkiG9w0BAQEFAASCAT4wggE6AgEAAkEAu1/a8f3G47cZ6pel
Hd3aONMNkGJ8vSCH5QjicuDm92VtVwKAACEjIoZSLYlJvPD+odL+lFzVQSmkneW7
VCGqYQIDAQABAkAcfNpg6GCQxoneLOghvlUrRotNZGvqpUOEAvHK3X7AJhz5SU4V
an36qvsAt5ghFMVM2iGvGaXbj0dAd+Jg64pxAiEA32Eh9mPtFmZhtIUMeGcPmPk
qIYCEuP8a/ZLmI9s4TsCIQDWvLQuvJsvfwPhi0TFAb5wqAET8X5LBFqtGX5QlUep
EwIgfNqM02Gc4wtLoqa2d4qPkYul3+uUW9hLd4Xsd6i/OS8CIQDT3elU+Rt+qIwW
u0cFrVvNYSV3HNzDfS9N/IoxTagfewIgpVxADe5c2EWbhCUkhN+ZCf38AKewK9TW
lQcDy4L+f14= -----END PRIVATE KEY-----",
    "public_certificate": "-----BEGIN CERTIFICATE-----
MIIBuzCCAwwGAWIBAgIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQAwHdENMAsgA1UE
AxMEVEVTVDELMAkGA1UEBhMCMVVMwHhcNMTgwNjA4MTgwOTAxWhcNMTkwNjA4MTgw
OTAxWjAcMQ0wCwYDVQQDEwRURVNUMQswCQYDVQQGEwJVUzBcMA0GCSqGSIb3DQEB
AQUAA0sAMEgCQQDaPvbqUJJFJ6NNTyK3Yb+ytSjJ9aa3yUmYTD9uMiP+6ycjxHWB
e8u9z6yCHsW03ync+dnhE5c5z8wuDAY0fv15AgMBAAGjgYowgYcwDAYDVR0TBAUw
AwEB/zALBgNVHQ8EBAMCAQYwHQYDVR0OBBYEFMJ7Ev/o/3+YNzYh5XNlqqjnw4zm
MEsGA1UdIwREMEKAFMJ7Ev/o/3+YNzYh5XNlqqjnw4zmoSCkhjAcMQ0wCwYDVQQD
EwRURVNUMQswCQYDVQQGEwJVU4IIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQADQQAv
DovYeyGNknjGI+TVNX6nDbyzf7zUPqnri0KuvObEeybrbPW45sgsnT5dyeE/32U
9Yr6lklklnkBTvBDTmLnrc -----END CERTIFICATE-----",
    "scope": "svm",
    "serial_number": "string",
    "svm": {
      "_links": {
```

```
    "self": {
      "href": "/api/resourcelink"
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "client",
  "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

security\_certificate

Name	Type	Description
_links	<a href="#">_links</a>	
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.

Name	Type	Description
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.
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.
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.



Name	Type	Description
type	string	Type of Certificate. The following types are supported: <ul style="list-style-type: none"> <li>• client - a certificate and its private key used by an SSL client in ONTAP.</li> <li>• server - a certificate and its private key used by an SSL server in ONTAP.</li> <li>• client_ca - a Certificate Authority certificate used by an SSL server in ONTAP to verify an SSL client certificate.</li> <li>• server_ca - a Certificate Authority certificate used by an SSL client in ONTAP to verify an SSL server certificate.</li> <li>• root_ca - a self-signed certificate used by ONTAP to sign other certificates by acting as a Certificate Authority.</li> <li>• enum: ["client", "server", "client_ca", "server_ca", "root_ca"]</li> </ul>
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code

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

## Create or install security certificates

POST /security/certificates

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.

### 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`

### Learn more

- [DOC /security/certificates](#)

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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.
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.

Name	Type	Description
svm	svm	SVM, applies only to SVM-scoped objects.
type	string	<p>Type of Certificate. The following types are supported:</p> <ul style="list-style-type: none"> <li>• client - a certificate and its private key used by an SSL client in ONTAP.</li> <li>• server - a certificate and its private key used by an SSL server in ONTAP.</li> <li>• client_ca - a Certificate Authority certificate used by an SSL server in ONTAP to verify an SSL client certificate.</li> <li>• server_ca - a Certificate Authority certificate used by an SSL client in ONTAP to verify an SSL server certificate.</li> <li>• root_ca - a self-signed certificate used by ONTAP to sign other certificates by acting as a Certificate Authority.</li> <li>• enum: ["client", "server", "client_ca", "server_ca", "root_ca"]</li> </ul>
uuid	string	Unique ID that identifies a certificate.

**Example request**



```

{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ca": "string",
  "common_name": "test.domain.com",
  "hash_function": "sha1",
  "intermediate_certificates": {
  },
  "private_key": "-----BEGIN PRIVATE KEY-----
MIIBVAIBADANBgkqhkiG9w0BAQEFAASCAT4wggE6AgEAAkEAu1/a8f3G47cZ6pe1
Hd3aONMNkGJ8vSCH5QjicuDm92VtVwkAACEjIoZSLYlJvPD+odL+lFzVQSmkneW7
VCGqYQIDAQABAkAcfNpg6GCQxoneLOghvlUrRotNZGvqpUOEAvHK3X7AJhZ5SU4V
an36qvsAt5ghFMVM2iGvGaXbj0dAd+Jg64pxAiEA32Eh9mPtFSmZhTIUMeGcPmPk
qIYCEuP8a/ZLmI9s4TsCIQDWvLQuvjSVfwPhi0TFab5wqAET8X5LBFqtGX5QlUep
EwIgfEnqM02Gc4wtLoqa2d4qPkYu13+uUW9hLd4Xsd6i/OS8CIQDT3elU+Rt+qIwW
u0cFrVvNYSV3HNzDfS9N/IoxTagfewIgpVxAdE5c2EWbhCUkhN+ZCf38AKewK9TW
lQcDy4L+f14= -----END PRIVATE KEY-----",
  "public_certificate": "-----BEGIN CERTIFICATE-----
MIIBuzCCAWWgAwIBAgIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQAWhDENMAsGA1UE
AxMEVEVETVDELMAkGA1UEBhMCMVVMWWhcNMTgwNjA4MTgwOTAxWhcNMTkwNjA4MTgw
OTAxWjAcMQ0wCwYDVQQDEwRURVNUMQswCQYDVQQGEwJVUzBcMA0GCSqGSIb3DQEB
AQUAA0sAMEgCQQDaPvbqUJJFJ6NNTyK3Yb+ytSjJ9aa3yUmYTD9uMiP+6ycjxHWB
e8u9z6yCHsW03ync+dnhE5c5z8wuDAY0fv15AgMBAAGjgYowgYcwDAYDVR0TBAUw
AwEB/zALBgNVHQ8EBAMCAQYwHQYDVR0OBByEFMJ7Ev/o/3+YNzYh5XNlqqjnw4zm
MEsGA1UdIwREMEKAFMJ7Ev/o/3+YNzYh5XNlqqjnw4zmoSCkHjAcMQ0wCwYDVQQD
EwRURVNUMQswCQYDVQQGEwJVU4IIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQADQQA
vDovYeyGNknjGI+TVNX6nDbyzf7zUPqnri0KuvObEeybrbPW45sgsnT5dyeE/32U
9Yr6lklknbTvbDTmLnrc -----END CERTIFICATE-----",
  "scope": "svm",
  "serial_number": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "client",
  "uuid": "string"
}

```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">security_certificate</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ca": "string",
    "common_name": "test.domain.com",
    "hash_function": "sha1",
    "intermediate_certificates": {
    },
    "private_key": "-----BEGIN PRIVATE KEY-----
MIIBVAIBADANBgkqhkiG9w0BAQEFAASCAT4wggE6AgEAAkEAu1/a8f3G47cZ6pel
Hd3aONMNkGJ8vSCH5QjicuDm92VtVwkaACEjIoZSLYlJvPD+odL+lFzVQSmkneW7
VCGqYQIDAQABAAkAcfNpg6GCQxoneLOghvlUrRotNZGvqpUOEAvHK3X7AJhz5SU4V
an36qvsAt5ghFMVM2iGvGaXbj0dAd+Jg64pxAiEA32Eh9mPtFsmZhTIUMeGcPmPk
qIYCEuP8a/ZLmI9s4TsCIQDWvLQuvJsvfwPhi0TFAb5wqAET8X5LBFqtGX5QlUep
EwIgfNqM02Gc4wtLoqa2d4qPkYul3+uUW9hLd4Xsd6i/OS8CIQDT3elU+Rt+qIwW
u0cFrVvNYSV3HNzDfS9N/IoxTagfewIgpVxADe5c2EWbhCUkhN+ZCf38AKewK9TW
lQcDy4L+f14= -----END PRIVATE KEY-----",
    "public_certificate": "-----BEGIN CERTIFICATE-----
MIIBuzCCAwwGAWIBAgIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQAwHdENMAsgA1UE
AxMEVEVTVDELMAkGA1UEBhMCMVVMwHhcNMTgwNjA4MTgwOTAxWhcNMTkwNjA4MTgw
OTAxWjAcMQ0wCwYDVQQDEwRURVNUMQswCQYDVQQGEwJVUzBcMA0GCSqGSIb3DQEB
AQUAA0sAMEgCQQDaPvbqUJJFJ6NNTyK3Yb+ytSjJ9aa3yUmYTD9uMiP+6ycjxHWB
e8u9z6yCHsW03ync+dnhE5c5z8wuDAY0fv15AgMBAAGjgYowgYcwDAYDVR0TBAUw
AwEB/zALBgNVHQ8EBAMCAQYwHQYDVR0OBBYEFMJ7Ev/o/3+YNzYh5XNlqqjnw4zm
MEsgA1UdIwREMEKAFMJ7Ev/o/3+YNzYh5XNlqqjnw4zmoSCkhJAcMQ0wCwYDVQQD
EwRURVNUMQswCQYDVQQGEwJVU4IIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQADQQAv
DovYeyGNknjGI+TVNX6nDbyzf7zUPqnri0KuvObEeybrbPW45sgsnT5dyeE/32U
9Yr6lklklnkBTvBDTmLnrc -----END CERTIFICATE-----",
    "scope": "svm",
    "serial_number": "string",
    "svm": {
      "_links": {
```



```

    "self": {
      "href": "/api/resourceLink"
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "client",
  "uuid": "string"
}

```

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
3735645	Cannot specify a value for serial. It is generated automatically
3735622	Certificate type not supported
3735664	Specified key size is not supported in FIPS mode
3735665	Specified hash function is not supported in FIPS mode
3735553	Failed to create self-signed Certificate
3735646	Failed to store the certificates
3735693	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	Certificate and private key do not match
52559973	The certificate has expired
52363366	Logic error: use of a dead object

Error Code	Description
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 invalid
3735632	Failed to extract Certificate Authority Information from the certificate

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

security\_certificate

Name	Type	Description
_links	<a href="#">_links</a>	
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.
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.
svm	svm	SVM, applies only to SVM-scoped objects.

Name	Type	Description
type	string	Type of Certificate. The following types are supported: <ul style="list-style-type: none"> <li>• client - a certificate and its private key used by an SSL client in ONTAP.</li> <li>• server - a certificate and its private key used by an SSL server in ONTAP.</li> <li>• client_ca - a Certificate Authority certificate used by an SSL server in ONTAP to verify an SSL client certificate.</li> <li>• server_ca - a Certificate Authority certificate used by an SSL client in ONTAP to verify an SSL server certificate.</li> <li>• root_ca - a self-signed certificate used by ONTAP to sign other certificates by acting as a Certificate Authority.</li> <li>• enum: ["client", "server", "client_ca", "server_ca", "root_ca"]</li> </ul>
uuid	string	Unique ID that identifies a certificate.

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
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

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

### 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

```
{
  "hash_function": "sha256",
  "signing_request": "'-----BEGIN CERTIFICATE REQUEST-----
MIICYDCCAUGCAQAwGzEMMAoGA1UEAxMDQUJDMQswCQYDVQQGEwJVUzCCASIwDQYJ
KoZIhvcNAQEBBQADggEPADCCAQoCggEBAPF+82SlqT3Vyu3Jx4IAwHcO5EGwLOxy
zQ6KNjz71Fca0n1/A1CbCPyOsSupGVObvdWxX7xLVMJ2SXB7h43GCqYyX6FXJO4F
HOpmLvB+jxdeiW7SdbiZyLUlsvA+oRO/uNlcug773QzdKLjJD64erZZMRUNbUJB8
bARxAUi0FPvgTraSQ0UW5sRLiGKeAyKA4wekYe1VgjHRTBizFbD4dI3njfva/2B1
jf+kkulgclJTtuJNtkgeimqMKyraYuleYcYk2K+C//0NuNOuPbDfTXCM7O61vik09
Szi8nLN7OXE9KoAA93U/BCpSfpl8XIb4cGnEr8hgVHOotZSo+KZBFxMCAwEAAaAA
MA0GCSqGSIB3DQEBCwUAA4IBAQC2vFYpvgsFrm5GnPx8tOBD1xsTyYjbWJMD8hAF
lFrvF9Sw9QGctDyacxkwgJhQx8l8JiIS5GOY6WwLb19FMkLQNAhDL9xF3WF7vfYq
RKgrz3bd/Vg96fsRZNYIPLGmoEaqLOh3FOGc2VbdsR9PwOn3fwthxkIRd6ds6/q
jc5cpSmVsCOgu+OKcpRXikyDbkWXfTZ1AhSfn6njBYfdZ9+PNAu/0JRQh5bX60nO
5heniTcaJLwUZP/CQ8nxHY0Wqy+1rAtM33d5cVmhu1BXQSIru/0ZkA/b9fK5Zv8E
ZMADYUoEvIG59VxhyCi8lzYf+Mxl8qBSF+Zdc4yWhzDqZtm9 -----END CERTIFICATE
REQUEST-----'"
}
```

### Response

Status: 200, Ok

Name	Type	Description
public_certificate	string	CA signed public key Certificate

#### Error

Status: Default

#### ONTAP Error Response Codes

Error Code	Description
3735628	Failed to use CA certificate for signing
3735665	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 invalid
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

Name	Type	Description
error	<a href="#">error</a>	



## Example error

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

## 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[ <a href="#">error_arguments</a> ]	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}
```

Deletes a security certificate.

### Related ONTAP commands

- `security certificate delete`

## Learn more

- [DOC /security/certificates](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

## 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	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	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}

Retrieves security certificates.

## Related ONTAP commands

- `security certificate show`

## Learn more

- [DOC /security/certificates](#)

## 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
<a href="#">_links</a>	<a href="#">_links</a>	
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.
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.
svm	svm	SVM, applies only to SVM-scoped objects.

Name	Type	Description
type	string	<p>Type of Certificate. The following types are supported:</p> <ul style="list-style-type: none"> <li>• client - a certificate and its private key used by an SSL client in ONTAP.</li> <li>• server - a certificate and its private key used by an SSL server in ONTAP.</li> <li>• client_ca - a Certificate Authority certificate used by an SSL server in ONTAP to verify an SSL client certificate.</li> <li>• server_ca - a Certificate Authority certificate used by an SSL client in ONTAP to verify an SSL server certificate.</li> <li>• root_ca - a self-signed certificate used by ONTAP to sign other certificates by acting as a Certificate Authority.</li> <li>• enum: ["client", "server", "client_ca", "server_ca", "root_ca"]</li> </ul>
uuid	string	Unique ID that identifies a certificate.

**Example response**

A large, empty rectangular box with a thin, dashed border, occupying most of the page. It is intended for an example response.



```

{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ca": "string",
  "common_name": "test.domain.com",
  "hash_function": "sha1",
  "intermediate_certificates": {
  },
  "private_key": "-----BEGIN PRIVATE KEY-----
MIIBVAIBADANBgkqhkiG9w0BAQEFAASCAT4wggE6AgEAAkEAu1/a8f3G47cZ6pe1
Hd3aONMNkGJ8vSCH5QjicuDm92VtVwkAACEjIoZSLYlJvPD+odL+lFzVQSmkneW7
VCGqYQIDAQABAAkAcfNpg6GCQxoneLOghvlUrRotNZGvqpUOEAvHK3X7AJhZ5SU4V
an36qvsAt5ghFMVM2iGvGaXbj0dAd+Jg64pxAiEA32Eh9mPtFSmZhTIUMeGcPmPk
qIYCEuP8a/ZLmI9s4TsCIQDWvLQuvjSVfwPhi0TFab5wqAET8X5LBFqtGX5QlUep
EwIgfEnqM02Gc4wtLoqa2d4qPkYu13+uUW9hLd4Xsd6i/OS8CIQDT3elU+Rt+qIwW
u0cFrVvNYSV3HNzDfS9N/IoxTagfewIgpVxAdE5c2EWbhCUkhN+ZCf38AKewK9TW
lQcDy4L+f14= -----END PRIVATE KEY-----",
  "public_certificate": "-----BEGIN CERTIFICATE-----
MIIBuzCCAWWgAwIBAgIIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQAWhDENMAsGA1UE
AxMEVEVETVDELMAkGA1UEBhMCMVVMWWhcNMTgwNjA4MTgwOTAxWhcNMTkwNjA4MTgw
OTAxWjAcMQ0wCwYDVQQDEwRURVNUMQswCQYDVQQGEwJVUzBcMA0GCSqGSIb3DQEB
AQUAA0sAMEgCQQDaPvbqUJJFJ6NNTyK3Yb+ytSjJ9aa3yUmYTD9uMiP+6ycjxHWB
e8u9z6yCHsW03ync+dnhE5c5z8wuDAY0fv15AgMBAAGjgYowgYcwDAYDVR0TBAUw
AwEB/zALBgNVHQ8EBAMCAQYwHQYDVR0OBBYEFMJ7Ev/o/3+YNzYh5XNlqqjnw4zm
MEsGA1UdIwREMEKAFMJ7Ev/o/3+YNzYh5XNlqqjnw4zmoSCkHjAcMQ0wCwYDVQQD
EwRURVNUMQswCQYDVQQGEwJVU4IIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQADQQA
vDovYeyGNknjGI+TVNX6nDbyzf7zUPqnri0KuvObEeybrbPW45sgsnT5dyeE/32U
9Yr6lklknbTvbDTmLnrC -----END CERTIFICATE-----",
  "scope": "svm",
  "serial_number": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "client",
  "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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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 onboard key management, the keys are stored in ONTAP in wrapped format using a key hierarchy created using the salted hash of the passphrase entered when configuring onboard key management. 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 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": "",
        "_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 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\" } }"
```

---

### 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 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\" }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "server": "10.225.89.34:5696",
      "_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": "",
      "_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}'
```

---

## 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

Retrieves key managers.

### Related ONTAP commands

- `security key-manager show-keystore`
- `security key-manager external show`

### Learn more

- [DOC /security/key-managers](#)

### Parameters

Name	Type	In	Required	Description
onboard.enabled	boolean	query	False	Filter by onboard.enabled
external.server_ca_certificates.uuid	string	query	False	Filter by external.server_ca_certificates.uuid
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.timeout	integer	query	False	Filter by external.servers.timeout

Name	Type	In	Required	Description
external.servers.user name	string	query	False	Filter by external.servers.use rname
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
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.
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.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">security_key_manager</a> ]	

## 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"
          }
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "server_ca_certificates": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "servers": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "server": "keyserver1.com:5698",
        "timeout": 60,
        "username": "username"
      }
    },
    "onboard": {
```

```

    "existing_passphrase": "The cluster password of length 32-256
ASCII characters.",
    "passphrase": "The cluster password of length 32-256 ASCII
characters."
  },
  "scope": "svm",
  "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

client\_certificate

Client certificate

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	Certificate UUID

server\_ca\_certificates

Security certificate object reference

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	Certificate UUID

key\_server\_readcreate

Name	Type	Description
_links	<a href="#">_links</a>	
server	string	External key server for key management. If no port is provided, a default port of 5696 is used.



Name	Type	Description
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	<a href="#">client_certificate</a>	Client certificate
server_ca_certificates	array[ <a href="#">server_ca_certificates</a> ]	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[ <a href="#">key_server_readcreate</a> ]	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.
passphrase	string	The cluster-wide passphrase. This is not audited.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### security\_key\_manager

Name	Type	Description
_links	<a href="#">_links</a>	
external	<a href="#">external</a>	Configures external key management
onboard	<a href="#">onboard</a>	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".
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

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

## Create a key manager

POST /security/key-managers

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.

### Related ONTAP commands

- `security key-manager external enable`
- `security key-manager onboard enable`

### Learn more

- [DOC /security/key-managers](#)

### Request Body

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>external</code>	<code>external</code>	Configures external key management
<code>onboard</code>	<code>onboard</code>	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
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	

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "external": {
    "client_certificate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "server_ca_certificates": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "servers": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "server": "keyserver1.com:5698",
      "timeout": 60,
      "username": "username"
    }
  },
  "onboard": {
    "existing_passphrase": "The cluster password of length 32-256 ASCII characters.",
    "passphrase": "The cluster password of length 32-256 ASCII characters."
  },
  "scope": "svm",
  "svm": {
    "_links": {
      "self": {
```

```
    "href": "/api/resourcelink"
  }
},
"name": "svm1",
"uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "string"
}
```

## Response

Status: 201, Created

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">security_key_manager</a> ]	

## 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"
          }
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "server_ca_certificates": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "servers": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "server": "keyserver1.com:5698",
        "timeout": 60,
        "username": "username"
      }
    },
    "onboard": {
```

```

    "existing_passphrase": "The cluster password of length 32-256
ASCII characters.",
    "passphrase": "The cluster password of length 32-256 ASCII
characters."
  },
  "scope": "svm",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourceLink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}

```

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
65536822	Multitenant key management is not supported in the current cluster version.
65536823	The SVM has key manager already configured.
65536878	External key management cannot be configured as one or more volume encryption keys of the SVM are stored in cluster key management server.
65536824	Multitenant key management is not supported in MetroCluster configurations.
65536038	A maximum of 4 active key servers are allowed.
65536876	External key management requires client and server CA certificates installed and with one or more key servers provided.
65536920	Onboard key manager passphrase length is incorrect.
65536871	Duplicate key management servers exist.
65536834	Failed to get existing key-server details for the SVM.



Error Code	Description
65536870	Key management servers already configured.
65536821	Certificate is not installed.
65536852	Failed to query supported KMIP protocol versions.
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.
65536916	Onboard key management is only supported for an admin SVM.
65536906	Onboard key management has already been configured at the partner site. Use the CLI to sync the onboard key management with the same passphrase.
65536907	Onboard key management is already configured. Use the CLI to sync any nodes with onboard key management configuration.
65536508	The platform does not support data at rest encryption.
65536310	Failed to setup onboard key management because the MetroCluster peer is unhealthy.
65536900	Onboard key management cannot be configured because this cluster is part of a MetroCluster configuration and the partner site has the external key manager configured.
65536903	Onboard key management has failed to configure on some nodes in the cluster. Use the CLI to sync the onboard key management configuration on failed nodes.
65536214	Failed to generate cluster key encryption key.
65536216	Failed to add cluster key encryption key.
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.
66060341	Failed to establish secure connection for a key management server due to network configuration issues.

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	<a href="#">href</a>	

client\_certificate

Client certificate

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	Certificate UUID

server\_ca\_certificates

Security certificate object reference

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	Certificate UUID

key\_server\_readcreate

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">client_certificate</a>	Client certificate
server_ca_certificates	array[ <a href="#">server_ca_certificates</a> ]	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[ <a href="#">key_server_readcreate</a> ]	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.
passphrase	string	The cluster-wide passphrase. This is not audited.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

security\_key\_manager

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
external	<a href="#">external</a>	Configures external key management
onboard	<a href="#">onboard</a>	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".
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}
```

Deletes a key manager.

### Related ONTAP commands

- `security key-manager external disable`
- `security key-manager onboard disable`

### Learn more

- [DOC /security/key-managers](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

### 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.
65536828	External key management is not enabled for the SVM.
65536242	One or more Storage Encryption devices are assigned an authentication key.
65536813	Encrypted kernel core files found.
65536817	Failed to determine if key manager is safe to disable.
65536827	Failed to determine if the SVM has any encrypted volumes.
65536867	Encrypted volumes are found for the SVM.
65536239	Encrypted volumes are found for the SVM.
196608301	Failed to determine the type of encryption.

Error Code	Description
196608305	NAE aggregates are found in the cluster.
65536242	One or more Storage Encryption devices are assigned an authentication key.
65536800	Failed to lookup onboard keys.
65536208	Failed to delete the SVM Key ID.
65536233	Internal error. Deletion of km_wrapped_kdb key database has failed for onboard key management.
65536234	Internal error. Deletion of cluster_kdb key database has failed for onboard key management.

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[ <a href="#">error_arguments</a> ]	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}

Retrieves key managers.

### Related ONTAP commands

- `security key-manager show-keystore`
- `security key-manager external show`

### Learn more

- [DOC /security/key-managers](#)

### 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
_links	<a href="#">_links</a>	
external	<a href="#">external</a>	Configures external key management
onboard	<a href="#">onboard</a>	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".
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "external": {
    "client_certificate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "server_ca_certificates": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "servers": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "server": "keyserver1.com:5698",
      "timeout": 60,
      "username": "username"
    }
  },
  "onboard": {
    "existing_passphrase": "The cluster password of length 32-256 ASCII characters.",
    "passphrase": "The cluster password of length 32-256 ASCII characters."
  },
  "scope": "svm",
  "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	<a href="#">href</a>	

client\_certificate

Client certificate

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	Certificate UUID

server\_ca\_certificates

Security certificate object reference

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	Certificate UUID

key\_server\_readcreate

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">client_certificate</a>	Client certificate
server_ca_certificates	array[ <a href="#">server_ca_certificates</a> ]	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[ <a href="#">key_server_readcreate</a> ]	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.
passphrase	string	The cluster-wide passphrase. This is not audited.

## svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

Updates a key manager.

### Related ONTAP commands

- `security key-manager external modify`
- `security key-manager onboard update-passphrase`

### Learn more

- [DOC /security/key-managers](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Key manager UUID

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
external	<a href="#">external</a>	Configures external key management

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".
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "external": {
    "client_certificate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "server_ca_certificates": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "servers": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "server": "keyserver1.com:5698",
      "timeout": 60,
      "username": "username"
    }
  },
  "onboard": {
    "existing_passphrase": "The cluster password of length 32-256 ASCII characters.",
    "passphrase": "The cluster password of length 32-256 ASCII characters."
  },
  "scope": "svm",
  "svm": {
    "_links": {
      "self": {
```



```

        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}

```

### 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.
65536828	External key management is not enabled for the SVM.
65536821	Certificate is not installed.
65536850	The 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.
65536150	New passphrase is same as old passphrase.
65536139	The existing passphrase value provided does not match the configured passphrase.
65536404	Passphrase does not match the accepted length.
65536802	Passphrase does not match the accepted length in common criteria mode.
65536408	Passphrase update failed on some nodes.
65536407	Passphrase update failed on some nodes.
65536406	Change of passphrase failed.

Error Code	Description
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.
66060341	Failed to establish secure connection for a key management server due to network configuration issues.

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	<a href="#">href</a>	

client\_certificate

Client certificate

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	Certificate UUID

server\_ca\_certificates

Security certificate object reference

Name	Type	Description
_links	<a href="#">_links</a>	
uuid	string	Certificate UUID

key\_server\_readcreate

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">client_certificate</a>	Client certificate
server_ca_certificates	array[ <a href="#">server_ca_certificates</a> ]	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[ <a href="#">key_server_readcreate</a> ]	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.
passphrase	string	The cluster-wide passphrase. This is not audited.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

security\_key\_manager

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
external	<a href="#">external</a>	Configures external key management
onboard	<a href="#">onboard</a>	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".
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

Retrieves key servers.

## Related ONTAP commands

- `security key-manager external show`

## Parameters

Name	Type	In	Required	Description
uuid	string	path	True	External key manager UUID
username	string	query	False	Filter by username
timeout	integer	query	False	Filter by timeout
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_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.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">key_server</a> ]	

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "password": "password",
    "records": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "password": "password",
      "server": "keyserver1.com:5698",
      "timeout": 60,
      "username": "username"
    },
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

records

Name	Type	Description
_links	<a href="#">_links</a>	
password	string	Password credentials for connecting with the key server. This is not audited.
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	

key\_server

Name	Type	Description
_links	<a href="#">_links</a>	
password	string	Password credentials for connecting with the key server. This is not audited.

Name	Type	Description
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 server is provided.
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.
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

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
<code>uuid</code>	string	path	True	External key manager UUID

## Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>password</code>	string	Password credentials for connecting with the key server. This is not audited.
<code>records</code>	array[ <a href="#">records</a> ]	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.
<code>server</code>	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.
<code>timeout</code>	integer	I/O timeout in seconds for communicating with the key server.
<code>username</code>	string	KMIP username credentials for connecting with the key server.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "password": "password",
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "password": "password",
    "server": "keyserver1.com:5698",
    "timeout": 60,
    "username": "username"
  },
  "server": "keyserver1.com:5698",
  "timeout": 60,
  "username": "username"
}
```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">key_server</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "password": "password",
    "records": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "password": "password",
      "server": "keyserver1.com:5698",
      "timeout": 60,
      "username": "username"
    },
    "server": "keyserver1.com:5698",
    "timeout": 60,
    "username": "username"
  }
}
```

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
65536822	Multitenant key management is not supported in the current cluster version.
65536828	External key management is not enabled for the SVM.
65536824	Multitenant key management is not supported in MetroCluster configurations.
65536038	A maximum of 4 active key servers are allowed.
65536871	Duplicate key management servers exist.
65536834	Failed to get existing key-server details for the SVM.
65536870	Key management servers already configured.
65536821	Certificate is not installed.
65536852	Failed to query supported KMIP protocol versions.
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.
66060341	Failed to establish secure connection for a key management server due to network configuration issues.

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	<a href="#">href</a>	

records

Name	Type	Description
_links	<a href="#">_links</a>	
password	string	Password credentials for connecting with the key server. This is not audited.
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	

key\_server

Name	Type	Description
_links	<a href="#">_links</a>	
password	string	Password credentials for connecting with the key server. This is not audited.
records	array[ <a href="#">records</a> ]	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.



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.
timeout	integer	I/O timeout in seconds for communicating with the key server.
username	string	KMIP username credentials for connecting with the key server.

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}

Deletes a key server.

### Related ONTAP commands

- `security key-manager external remove-servers`

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	
server	string	path	True	

### 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.
65536828	External key management is not enabled for the SVM.
65536824	Multitenant key management is not supported in MetroCluster configurations.
65536843	Key management server is not configured for the SVM.
65536700	The key server contains keys that are currently in use and not available from any other configured key server in the SVM.

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	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}

Retrieves key servers configured in an external key manager.

## Related ONTAP commands

- `security key-manager external show`

## 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
<a href="#">_links</a>	<a href="#">_links</a>	
password	string	Password credentials for connecting with the key server. This is not audited.
records	array[ <a href="#">records</a> ]	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.
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.

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "password": "password",
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "password": "password",
    "server": "keyserver1.com:5698",
    "timeout": 60,
    "username": "username"
  },
  "server": "keyserver1.com:5698",
  "timeout": 60,
  "username": "username"
}
```

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

records

Name	Type	Description
_links	<a href="#">_links</a>	
password	string	Password credentials for connecting with the key server. This is not audited.
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	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code

Name	Type	Description
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}

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
_links	<a href="#">_links</a>	
password	string	Password credentials for connecting with the key server. This is not audited.
records	array[ <a href="#">records</a> ]	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.



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 <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.

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "password": "password",
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "password": "password",
  "server": "keyserver1.com:5698",
  "timeout": 60,
  "username": "username"
},
"server": "keyserver1.com:5698",
"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.
65536828	External key management is not enabled for the SVM.
65536824	Multitenant key management is not supported in MetroCluster configurations.
65536843	Key management server is not configured for the SVM.
65536846	Missing password.
65536845	Missing username.

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	<a href="#">href</a>	

records

Name	Type	Description
_links	<a href="#">_links</a>	
password	string	Password credentials for connecting with the key server. This is not audited.
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	

key\_server

Name	Type	Description
_links	<a href="#">_links</a>	
password	string	Password credentials for connecting with the key server. This is not audited.
records	array[ <a href="#">records</a> ]	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.

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.
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[ <a href="#">error_arguments</a> ]	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 operation 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: application/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: application/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

Retrieves the login banner and messages of the day (MOTD) configured in the cluster and in specific SVMs.

### Learn more

- [DOC /security/login/messages](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	query	False	Filter by uuid
show_cluster_message	boolean	query	False	Filter by show_cluster_message
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
scope	string	query	False	Filter by scope
banner	string	query	False	Filter by banner
message	string	query	False	Filter by 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	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.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">login_messages</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "scope": "svm",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	

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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Retrieve login messages configuration by UUID

GET /security/login/messages/{uuid}

Retrieves the login messages configuration by UUID.

#### Learn more

- [DOC /security/login/messages](#)

## 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
<a href="#">_links</a>	<a href="#">_links</a>	
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.

Name	Type	Description
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. New lines are supplied as either LF or CRLF but are always returned as LF. Optional in the PATCH body.</p>
scope	string	<p>Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".</p>
show_cluster_message	boolean	<p>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.</p>
svm	svm	<p>SVM, applies only to SVM-scoped objects.</p>
uuid	string	<p>The unique identifier (ID) of the login messages configuration.</p>

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "scope": "svm",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}
```

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.



## Update login messages configuration

PATCH /security/login/messages/{uuid}

Updates the login messages configuration.  
There are no required fields. An empty body will make no modifications.

### Learn more

- [DOC /security/login/messages](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Login messages configuration UUID

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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.

Name	Type	Description
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. New lines are supplied as either LF or CRLF but are always returned as LF. Optional in the PATCH body.</p>
scope	string	<p>Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".</p>
show_cluster_message	boolean	<p>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.</p>
svm	svm	<p>SVM, applies only to SVM-scoped objects.</p>
uuid	string	<p>The unique identifier (ID) of the login messages configuration.</p>

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "scope": "svm",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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 show_cluster_message property.

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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
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.

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	SVM, applies only to SVM-scoped objects.
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[ <a href="#">error_arguments</a> ]	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 operations 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 operation is allowed on APIs with *api/cluster* as the prefix; while POST, PATCH and DELETE operations 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 `<em>/api/security/roles</em>`, 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 `<em>/api/security/roles</em>` or its derivatives. Legacy roles are managed using the CLI commands "security login role &lt;create \ | modify \ | delete&gt; -role <rolename>".</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 -k -u <cluster-admin>:<password> -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/applicat
ion/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 GET performed on the */api/svm/svms* API.



```

# The API:
POST "/api/security/roles"

# The call:
curl -k -u <cluster-admin>:<password> -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/applicat
ion/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 -k -u <cluster_admin>:<password> -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 -k -u <cluster_admin>:<password> -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 -k -u <cluster_admin>:<password> -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 -k -u <cluster_admin>:<password> -X GET "https://<mgmt-
ip>/api/security/roles/?name=vsadmin*"

# Scoped call to retrieve the predefined roles:
curl -k -u <cluster_admin>:<password> -X GET "https://<mgmt-
ip>/api/security/roles/?builtin=true"

# Scoped call to retrieve the custom roles:
curl -k -u <cluster_admin>:<password> -X GET "https://<mgmt-
ip>/api/security/roles/?builtin=false"
```

### Retrieve a list of roles configured in the cluster

GET /security/roles

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
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">role</a> ]	

## 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": "cluster"
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

owner

Owner name and UUID that uniquely identifies the role.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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
<a href="#">_links</a>	<a href="#">_links</a>	
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
<a href="#">_links</a>	<a href="#">_links</a>	
builtin	boolean	Indicates if this is a built-in (pre-defined) role which cannot be modified or deleted.
name	string	Role name
owner	<a href="#">owner</a>	Owner name and UUID that uniquely identifies the role.
privileges	array[ <a href="#">role_privilege</a> ]	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[ <a href="#">error_arguments</a> ]	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

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](#)

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>builtin</code>	boolean	Indicates if this is a built-in (pre-defined) role which cannot be modified or deleted.
<code>name</code>	string	Role name
<code>owner</code>	<a href="#">owner</a>	Owner name and UUID that uniquely identifies the role.
<code>privileges</code>	array[ <a href="#">role_privilege</a> ]	The list of privileges that this role has been granted.
<code>scope</code>	string	Scope of the entity. set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

## Example request

```
{
  "_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": "cluster"
}
```

## Response

Status: 201, Created

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
13434891	UUID lookup failed for Vserver roles.

Error Code	Description
13434890	Vserver-Id failed for Vserver roles.
13434892	Roles is a required field.
13434893	SVM does not exist.
5636169	Invalid character in URI.
5636170	URI does not exist.
5636129	Role with given name has not been defined.
5636144	Invalid value specified for access level.
5636171	Role already exists in legacy role table.
5636143	A Vserver admin cannot use the API with this access level.

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	<a href="#">href</a>	

owner

Owner name and UUID that uniquely identifies the role.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
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	<a href="#">_links</a>	
builtin	boolean	Indicates if this is a built-in (pre-defined) role which cannot be modified or deleted.
name	string	Role name

Name	Type	Description
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.

## 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 -k -u <cluster_admin>:<password> -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": "svm1",
    "_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 -k -u <cluster_admin>:<password> -X DELETE "https://<mgmt-
ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_role1"
```

## Delete a role

```
DELETE /security/roles/{owner.uuid}/{name}
```

Delete 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	
name	string	path	True	

### Response

```
Status: 200, Ok
```

### Error

```
Status: Default
```

### ONTAP Error Response Codes

Error Code	Description
5636173	Features requires an effective cluster version of 9.6 or later.
5636172	User accounts detected with this role assigned. Modify/delete those accounts before deleting this role.
13434893	SVM does not exist.
13434890	Vserver-Id failed for Vserver roles.

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[ <a href="#">error_arguments</a> ]	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}

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	
name	string	path	True	

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">role</a> ]	

## 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": "cluster"
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

owner

Owner name and UUID that uniquely identifies the role.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
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
<a href="#">_links</a>	<a href="#">_links</a>	
builtin	boolean	Indicates if this is a built-in (pre-defined) role which cannot be modified or deleted.
name	string	Role name
owner	<a href="#">owner</a>	Owner name and UUID that uniquely identifies the role.
privileges	array[ <a href="#">role_privilege</a> ]	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[ <a href="#">error_arguments</a> ]	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 call performed on one of the following APIs:  
<em>/api/security/roles</em> for all the roles  
<em>/api/security/roles/?scope=svm</em> for SVM-scoped roles  
<em>/api/security/roles/?owner.name=<svm-name>&lt;i>&lt;/i></em> for roles in a specific SVM This API response contains the complete URI for each role and can be used after suffixing it with `_"privileges"`.</em>.</svm-name>\_



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 -k -u <cluster_admin>:<password> -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 -k -u <cluster_admin>:<password> -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
```

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	
name	string	path	True	

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">role_privilege</a> ]	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

role\_privilege

A tuple containing the REST endpoint and the access level assigned to that endpoint.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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

Add 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").

**Optional parameters**

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	Owner UUID of the role.
name	string	path	True	Role name

**Request Body**

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
access	string	Access level for the REST endpoint.
path	string	REST URI/endpoint

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "access": "readonly",
  "path": "/api/storage/volumes"
}
```

### Response

Status: 201, Created

### Error

Status: Default

### ONTAP Error Response Codes

Error Code	Description
13434891	UUID LookUp failed for Vserver roles.
13434890	Vserver-Id failed for Vserver roles.
13434892	Roles is a required field.
13434893	SVM does not exist.
5636173	This feature requires an effective cluster version of 9.6 or later.
5636129	Role with given name has not been defined.
5636169	Invalid character in URI.
5636170	URI does not exist.



Error Code	Description
5636175	Vserver admin cannot have access to given API.
5636144	Invalid value specified for the access level.
5636143	A Vserver admin cannot use the API with this access level.

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	<a href="#">href</a>	

role\_privilege

A tuple containing the REST endpoint and the access level assigned to that endpoint.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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 call performed on one of the following APIs:<br> <em>/api/security/roles</em> for all roles<br> <em>/api/security/roles/?scope=svm</em> for SVM-scoped roles<br> <em>/api/security/roles/?owner.name=<svm-name>&lt;i>/i>&gt;</em> 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></em>



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 -k -u <cluster_admin>:<password> -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 -k -u <cluster_admin>:<password> -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 -k -u <cluster_admin>:<password> -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}
```

Delete 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	
name	string	path	True	
path	string	path	True	

### Response

```
Status: 200, Ok
```

### Error

```
Status: Default
```

### ONTAP Error Response Codes

Error Code	Description
5636173	Features requires an effective cluster version of 9.6 or later.
5636172	User accounts detected with this role assigned. Modify/delete those accounts before deleting this role.
13434893	SVM does not exist.
13434890	Vserver-Id failed for Vserver roles.

Name	Type	Description
error	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	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}

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	
name	string	path	True	
path	string	path	True	

## Response

Status: 200, Ok

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>access</code>	string	Access level for the REST endpoint.
<code>path</code>	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	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}
```

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	Owner UUID of the role.
name	string	path	True	Role name
path	string	path	True	REST API path

## Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
access	string	Access level for the REST endpoint.
path	string	REST URI/endpoint

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "access": "readonly",
  "path": "/api/storage/volumes"
}
```

## 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

role\_privilege

A tuple containing the REST endpoint and the access level assigned to that endpoint.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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 "mirror-vault", "sync-mirror" and "strict-sync-mirror". When applied to a SnapMirror relationship, the SnapMirror policy controls the behavior of the relationship and specifies the configuration attributes for that relationship.

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

### Retrieve SnapMirror async and sync policy types

GET /snapmirror/policies

Retrieves SnapMirror policies of type "mirror-vault", "sync-mirror" and "strict-sync-mirror".

## Related ONTAP commands

- `snapmirror policy show`

## Example

The following example shows how to retrieve a collection of SnapMirror policies.

```
GET "/api/storage/snapmirror/policies"
```

## Learn more

- [DOC /snapmirror/policies](#)

## Parameters

Name	Type	In	Required	Description
uuid	string	query	False	Filter by uuid
comment	string	query	False	Filter by comment
transfer_schedule.uuid	string	query	False	Filter by transfer_schedule.uuid
transfer_schedule.name	string	query	False	Filter by transfer_schedule.name
network_compression_enabled	boolean	query	False	Filter by network_compression_enabled
sync_type	string	query	False	Filter by sync_type
sync_common_snapshot_schedule.uuid	string	query	False	Filter by sync_common_snapshot_schedule.uuid
sync_common_snapshot_schedule.name	string	query	False	Filter by sync_common_snapshot_schedule.name
retention.prefix	string	query	False	Filter by retention.prefix
retention.count	integer	query	False	Filter by retention.count

Name	Type	In	Required	Description
retention.label	string	query	False	Filter by retention.label
retention.creation_schedule.uuid	string	query	False	Filter by retention.creation_schedule.uuid
retention.creation_schedule.name	string	query	False	Filter by retention.creation_schedule.name
throttle	integer	query	False	Filter by throttle
identity_preservation	string	query	False	Filter by identity_preservation
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
type	string	query	False	Filter by 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	The default is true for GET calls. When set to false, only the number of records is returned.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">snapmirror_policy</a> ]	



## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "comment": "string",
    "identity_preservation": "full",
    "name": "MirrorAndVault",
    "retention": {
      "count": 7,
      "creation_schedule": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "label": "hourly",
      "prefix": "string"
    },
    "scope": "svm",
    "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": "sync",
"throttle": 0,
"transfer_schedule": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "weekly",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"type": "async",
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
}

```

## Error

Status: Default

## ONTAP Error Response codes

Error code	Description
13303842	Invalid SnapMirror policy type.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

creation\_schedule

Schedule used to create Snapshot copies on the destination for long term retention.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">creation_schedule</a>	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

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

transfer\_schedule

The schedule used to update asynchronous relationships.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

snapmirror\_policy

SnapMirror policy information

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	

Name	Type	Description
comment	string	Comment associated with the policy.
identity_preservation	string	Specifies which configuration of the source SVM is replicated to destination. This field is applicable only for SVM data protection and async policies.
name	string	
network_compression_enabled	boolean	Specifies whether network compression is enabled for transfers. This is applicable only to async policies.
retention	array[ <a href="#">snapmirror_policy_rule</a> ]	Policy on Snapshot copy retention. This is applicable only to async policies.
scope	string	Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
sync_common_snapshot_schedule	<a href="#">sync_common_snapshot_schedule</a>	Schedule used to create common Snapshot copies for synchronous relationships.
sync_type	string	
throttle	integer	Throttle in KB/s. Default to unlimited.
transfer_schedule	<a href="#">transfer_schedule</a>	The schedule used to update asynchronous relationships.
type	string	
uuid	string	

#### error\_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

Creates a SnapMirror policy. The parameter "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", "retention" and "transfer\_schedule" 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"}'
```

## Learn more

- [DOC /snapmirror/policies](#)



## Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	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.

## Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	Comment associated with the policy.
identity_preservation	string	Specifies which configuration of the source SVM is replicated to destination. This field is applicable only for SVM data protection and async policies.
name	string	

Name	Type	Description
network_compression_enabled	boolean	Specifies whether network compression is enabled for transfers. This is applicable only to async policies.
retention	array[ <a href="#">snapmirror_policy_rule</a> ]	Policy on Snapshot copy retention. This is applicable only to async policies.
scope	string	Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
sync_common_snapshot_schedule	<a href="#">sync_common_snapshot_schedule</a>	Schedule used to create common Snapshot copies for synchronous relationships.
sync_type	string	
throttle	integer	Throttle in KB/s. Default to unlimited.
transfer_schedule	<a href="#">transfer_schedule</a>	The schedule used to update asynchronous relationships.
type	string	
uuid	string	

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "identity_preservation": "full",
  "name": "MirrorAndVault",
  "retention": {
    "count": 7,
    "creation_schedule": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "weekly",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "label": "hourly",
    "prefix": "string"
  },
  "scope": "svm",
  "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": "sync",
  "throttle": 0,
}
```

```
"transfer_schedule": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "weekly",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"type": "async",
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

### Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

### Example response

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

### Error

Status: Default

### ONTAP Error Response codes

Error code	Description
13303850	Invalid input parameter

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	<a href="#">href</a>	

creation\_schedule

Schedule used to create Snapshot copies on the destination for long term retention.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">creation_schedule</a>	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

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

transfer\_schedule

The schedule used to update asynchronous relationships.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

snapmirror\_policy

SnapMirror policy information

Name	Type	Description
_links	<a href="#">_links</a>	
comment	string	Comment associated with the policy.
identity_preservation	string	Specifies which configuration of the source SVM is replicated to destination. This field is applicable only for SVM data protection and async policies.
name	string	

Name	Type	Description
network_compression_enabled	boolean	Specifies whether network compression is enabled for transfers. This is applicable only to async policies.
retention	array[ <a href="#">snapmirror_policy_rule</a> ]	Policy on Snapshot copy retention. This is applicable only to async policies.
scope	string	Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
sync_common_snapshot_schedule	<a href="#">sync_common_snapshot_schedule</a>	Schedule used to create common Snapshot copies for synchronous relationships.
sync_type	string	
throttle	integer	Throttle in KB/s. Default to unlimited.
transfer_schedule	<a href="#">transfer_schedule</a>	The schedule used to update asynchronous relationships.
type	string	
uuid	string	

#### job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}
```

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	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.

### Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	Comment associated with the policy.
identity_preservation	string	Specifies which configuration of the source SVM is replicated to destination. This field is applicable only for SVM data protection and async policies.
name	string	
network_compression_enabled	boolean	Specifies whether network compression is enabled for transfers. This is applicable only to async policies.
retention	array[ <a href="#">snapmirror_policy_rule</a> ]	Policy on Snapshot copy retention. This is applicable only to async policies.

Name	Type	Description
scope	string	Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
svm	svm	SVM, applies only to SVM-scoped objects.
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": "full",
  "name": "MirrorAndVault",
  "retention": {
    "count": 7,
    "creation_schedule": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "weekly",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "label": "hourly",
    "prefix": "string"
  },
  "scope": "svm",
  "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": "sync",
  "throttle": 0,
}
```

```

"transfer_schedule": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "weekly",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"type": "async",
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}

```

## Error

Status: Default

## ONTAP Error Response codes

Error code	Description
13303842	Invalid SnapMirror policy type.

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	<a href="#">href</a>	

creation\_schedule

Schedule used to create Snapshot copies on the destination for long term retention.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">creation_schedule</a>	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

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

#### transfer\_schedule

The schedule used to update asynchronous relationships.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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

## Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	Comment associated with the policy.
identity_preservation	string	Specifies which configuration of the source SVM is replicated to destination. This field is applicable only for SVM data protection and async policies.
name	string	
network_compression_enabled	boolean	Specifies whether network compression is enabled for transfers. This is applicable only to async policies.
retention	array[ <a href="#">snapmirror_policy_rule</a> ]	Policy on Snapshot copy retention. This is applicable only to async policies.
scope	string	Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
sync_common_snapshot_schedule	<a href="#">sync_common_snapshot_schedule</a>	Schedule used to create common Snapshot copies for synchronous relationships.
sync_type	string	
throttle	integer	Throttle in KB/s. Default to unlimited.
transfer_schedule	<a href="#">transfer_schedule</a>	The schedule used to update asynchronous relationships.
type	string	

Name	Type	Description
uuid	string	

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "identity_preservation": "full",
  "name": "MirrorAndVault",
  "retention": {
    "count": 7,
    "creation_schedule": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "weekly",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "label": "hourly",
    "prefix": "string"
  },
  "scope": "svm",
  "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": "sync",
  "throttle": 0,
}
```

```
"transfer_schedule": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "weekly",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"type": "async",
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

### Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

### Example response

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

### Error

Status: Default

### ONTAP Error Response codes

Error code	Description
13303842	Invalid SnapMirror policy type
13303843	Conflicting values between SnapMirror policy and SnapMirror relationships for either 'transfer_schedule, throttle or identity_preservation' properties
13303850	Invalid input parameter

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	<a href="#">href</a>	

creation\_schedule

Schedule used to create Snapshot copies on the destination for long term retention.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">creation_schedule</a>	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

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

#### transfer\_schedule

The schedule used to update asynchronous relationships.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Job schedule name
uuid	string	Job schedule UUID

#### snapmirror\_policy

SnapMirror policy information

Name	Type	Description
_links	<a href="#">_links</a>	
comment	string	Comment associated with the policy.
identity_preservation	string	Specifies which configuration of the source SVM is replicated to destination. This field is applicable only for SVM data protection and async policies.
name	string	

Name	Type	Description
network_compression_enabled	boolean	Specifies whether network compression is enabled for transfers. This is applicable only to async policies.
retention	array[ <a href="#">snapmirror_policy_rule</a> ]	Policy on Snapshot copy retention. This is applicable only to async policies.
scope	string	Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
sync_common_snapshot_schedule	<a href="#">sync_common_snapshot_schedule</a>	Schedule used to create common Snapshot copies for synchronous relationships.
sync_type	string	
throttle	integer	Throttle in KB/s. Default to unlimited.
transfer_schedule	<a href="#">transfer_schedule</a>	The schedule used to update asynchronous relationships.
type	string	
uuid	string	

#### job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
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 create a synchronous relationship between FlexVol volumes, which provides zero RPO data protection. It supports the SnapMirror policy types "mirror-vault", "sync-mirror", and "strict-sync-mirror". You can create a relationship between the source and destination which can be used to 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.

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

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.
healthy	boolean	query	False	Filter by healthy
source.svm.uuid	string	query	False	Filter by source.svm.uuid
source.svm.name	string	query	False	Filter by source.svm.name
source.path	string	query	False	Filter by source.path
lag_time	string	query	False	Filter by lag_time
exported_snapshot	string	query	False	Filter by exported_snapshot
uuid	string	query	False	Filter by uuid
restore	boolean	query	False	Filter by restore
destination.svm.uuid	string	query	False	Filter by destination.svm.uuid

Name	Type	In	Required	Description
destination.svm.name	string	query	False	Filter by destination.svm.name
destination.path	string	query	False	Filter by destination.path
transfer.bytes_transferred	integer	query	False	Filter by transfer.bytes_transferred
transfer.uuid	string	query	False	Filter by transfer.uuid
transfer.state	string	query	False	Filter by transfer.state
policy.name	string	query	False	Filter by policy.name
policy.type	string	query	False	Filter by policy.type
policy.uuid	string	query	False	Filter by policy.uuid
state	string	query	False	Filter by state
unhealthy_reason.code	integer	query	False	Filter by unhealthy_reason.code
unhealthy_reason.message	string	query	False	Filter by unhealthy_reason.message
unhealthy_reason.parameters	string	query	False	Filter by unhealthy_reason.parameters
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">snapmirror_relationship</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "destination": {
      "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": "MirrorAndVault",
      "type": "async",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    },
    "source": {
      "path": "svm1:volume1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  }
}
```



```

    }
  },
  "state": "broken_off",
  "transfer": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "state": "aborted",
  "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror\_endpoint

Endpoint of a SnapMirror relationship.

Name	Type	Description
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: <ul style="list-style-type: none"><li>• example: svm1:volume1</li><li>• readCreate: 1</li></ul>
svm	<a href="#">svm</a>	

policy

Basic policy information of the relationship.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
type	string	
uuid	string	

transfer

Basic information on the current transfer.

Name	Type	Description
_links	<a href="#">_links</a>	
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

Name	Type	Description
destination	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.
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.

Name	Type	Description
policy	<a href="#">policy</a>	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 field is applicable only for relationships with FlexGroup or FlexVol endpoints and when the PATCH state is "snapmirrored".
quick_resync	boolean	Set to true to reduce resync time by not preserving storage efficiency. This field is applicable only for relationships with FlexVol endpoints and when the PATCH state is "snapmirrored".
recover_after_break	boolean	Set to true to recover from a failed SnapMirror break operation on a FlexGroup relationship. This restores all destination FlexGroup constituents to the latest Snapshot copy, and any writes to the read-write constituents are lost. This field is applicable only for SnapMirror relationships with FlexGroup endpoints and when the PATCH state is "broken_off".
restore	boolean	Set to true to create a relationship for restore. To trigger restore-transfer, use transfers POST on the restore relationship.
restore_to_snapshot	string	Specifies the Snapshot copy to restore to on the destination after a break operation. This field is applicable only for SnapMirror relationships with FlexVol endpoints and when the PATCH state is "broken_off".
source	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.

Name	Type	Description
state	string	State of the relationship. To initialize the relationship, PATCH the state to "snapmirrored". To break the relationship, PATCH the state to "broken_off". To resync the broken relationship, PATCH the state to "snapmirrored" for relationships with async policy type or "in_sync" for relationships with sync policy type. To pause the relationship, suspending further transfers, PATCH the state to "paused". To resume transfers for a paused relationship, PATCH the state to "snapmirrored" or "in_sync". The entries "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships of the sync policy type. 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" fields using GET to determine if the transition is successful.
transfer	<a href="#">transfer</a>	Basic information on the current transfer.
unhealthy_reason	array[ <a href="#">snapmirror_error</a> ]	Reason the relationship is not healthy. It is a concatenation of up to four levels of error messages.
uuid	string	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create a SnapMirror relationship

POST /snapmirror/relationships

Creates a SnapMirror relationship. This API must be executed on the cluster containing the destination endpoint.

### Required properties

- `source.path` - Path to the source endpoint of the SnapMirror relationship.
- `destination.path` - Path to the destination endpoint of the SnapMirror relationship.

### Recommended optional properties

- `policy.name` or `policy.uuid` - Policy governing the SnapMirror relationship.

### Default property values

If not specified in POST, the following default property values are assigned:

- `policy.name` - *MirrorAndVault*
- `restore` - *false*

### Related ONTAP commands

- `snapmirror create`

### Examples

The following examples show how to create FlexVol, FlexGroup and SVM 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":
"test_vserv_src:src_vol_rw"}, "destination": { "path":
"test_vserv_dst:dst_vol_rw"}}'
```

Creating a FlexGroup SnapMirror relationship of type XDP.

```
POST "/api/snapmirror/relationships/" '{"source": {"path":  
"test_vserv_src:source_flexgrp"}, "destination": { "path":  
"test_vserv_dst: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":  
"test_vserv_src:src_vol_rw"}, "destination": { "path":  
"test_vserv_dst:dst_vol_rw"}, "restore": "true"}'
```

**Learn more**

- [DOC /snapmirror/relationships](#)

**Parameters**



Name	Type	In	Required	Description
return_timeout	integer	query	False	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.

#### Request Body

Name	Type	Description
destination	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.
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	<a href="#">policy</a>	Basic policy information of the relationship.

Name	Type	Description
preserve	boolean	Set to true on resync to preserve Snapshot copies on the destination that are newer than the latest common Snapshot copy. This field is applicable only for relationships with FlexGroup or FlexVol endpoints and when the PATCH state is "snapmirrored".
quick_resync	boolean	Set to true to reduce resync time by not preserving storage efficiency. This field is applicable only for relationships with FlexVol endpoints and when the PATCH state is "snapmirrored".
recover_after_break	boolean	Set to true to recover from a failed SnapMirror break operation on a FlexGroup relationship. This restores all destination FlexGroup constituents to the latest Snapshot copy, and any writes to the read-write constituents are lost. This field is applicable only for SnapMirror relationships with FlexGroup endpoints and when the PATCH state is "broken_off".
restore	boolean	Set to true to create a relationship for restore. To trigger restore-transfer, use transfers POST on the restore relationship.
restore_to_snapshot	string	Specifies the Snapshot copy to restore to on the destination after a break operation. This field is applicable only for SnapMirror relationships with FlexVol endpoints and when the PATCH state is "broken_off".
source	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.

Name	Type	Description
state	string	<p>State of the relationship. To initialize the relationship, PATCH the state to "snapmirrored". To break the relationship, PATCH the state to "broken_off". To resync the broken relationship, PATCH the state to "snapmirrored" for relationships with async policy type or "in_sync" for relationships with sync policy type. To pause the relationship, suspending further transfers, PATCH the state to "paused". To resume transfers for a paused relationship, PATCH the state to "snapmirrored" or "in_sync". The entries "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships of the sync policy type. 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" fields using GET to determine if the transition is successful.</p>
transfer	<a href="#">transfer</a>	Basic information on the current transfer.
unhealthy_reason	array[ <a href="#">snapmirror_error</a> ]	Reason the relationship is not healthy. It is a concatenation of up to four levels of error messages.
uuid	string	

## Example request

```
{
  "destination": {
    "path": "svm1:volumel",
    "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": "MirrorAndVault",
    "type": "async",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "source": {
    "path": "svm1:volumel",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "state": "broken_off",
  "transfer": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
}
```

```

    "state": "aborted",
    "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": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}

```

#### Error

Status: Default

## ONTAP Error Response codes

Error code	Description
13303841	Cross platform relationships are not allowed
13303821	Invalid SnapMirror policy UUID
13303819	Could not retrieve SnapMirror policy information
13303852	destination.path provided does not contain "\:"
13303853	Restore relationships are not supported for SVM-DR endpoints

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	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror\_endpoint

Endpoint of a SnapMirror relationship.

Name	Type	Description
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: <ul style="list-style-type: none"><li>• example: svm1:volume1</li><li>• readCreate: 1</li></ul>
svm	<a href="#">svm</a>	

policy

Basic policy information of the relationship.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
type	string	
uuid	string	

transfer

Basic information on the current transfer.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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

Name	Type	Description
destination	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.
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	<a href="#">policy</a>	Basic policy information of the relationship.



Name	Type	Description
preserve	boolean	Set to true on resync to preserve Snapshot copies on the destination that are newer than the latest common Snapshot copy. This field is applicable only for relationships with FlexGroup or FlexVol endpoints and when the PATCH state is "snapmirrored".
quick_resync	boolean	Set to true to reduce resync time by not preserving storage efficiency. This field is applicable only for relationships with FlexVol endpoints and when the PATCH state is "snapmirrored".
recover_after_break	boolean	Set to true to recover from a failed SnapMirror break operation on a FlexGroup relationship. This restores all destination FlexGroup constituents to the latest Snapshot copy, and any writes to the read-write constituents are lost. This field is applicable only for SnapMirror relationships with FlexGroup endpoints and when the PATCH state is "broken_off".
restore	boolean	Set to true to create a relationship for restore. To trigger restore-transfer, use transfers POST on the restore relationship.
restore_to_snapshot	string	Specifies the Snapshot copy to restore to on the destination after a break operation. This field is applicable only for SnapMirror relationships with FlexVol endpoints and when the PATCH state is "broken_off".
source	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.

Name	Type	Description
state	string	State of the relationship. To initialize the relationship, PATCH the state to "snapmirrored". To break the relationship, PATCH the state to "broken_off". To resync the broken relationship, PATCH the state to "snapmirrored" for relationships with async policy type or "in_sync" for relationships with sync policy type. To pause the relationship, suspending further transfers, PATCH the state to "paused". To resume transfers for a paused relationship, PATCH the state to "snapmirrored" or "in_sync". The entries "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships of the sync policy type. 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" fields using GET to determine if the transition is successful.
transfer	<a href="#">transfer</a>	Basic information on the current transfer.
unhealthy_reason	array[ <a href="#">snapmirror_error</a> ]	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
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Delete a SnapMirror relationship

```
DELETE /snapmirror/relationships/{uuid}
```

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.

#### 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.
return_timeout	integer	query	False	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.

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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

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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}



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
destination	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.
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.

Name	Type	Description
policy	<a href="#">policy</a>	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 field is applicable only for relationships with FlexGroup or FlexVol endpoints and when the PATCH state is "snapmirrored".
quick_resync	boolean	Set to true to reduce resync time by not preserving storage efficiency. This field is applicable only for relationships with FlexVol endpoints and when the PATCH state is "snapmirrored".
recover_after_break	boolean	Set to true to recover from a failed SnapMirror break operation on a FlexGroup relationship. This restores all destination FlexGroup constituents to the latest Snapshot copy, and any writes to the read-write constituents are lost. This field is applicable only for SnapMirror relationships with FlexGroup endpoints and when the PATCH state is "broken_off".
restore	boolean	Set to true to create a relationship for restore. To trigger restore-transfer, use transfers POST on the restore relationship.
restore_to_snapshot	string	Specifies the Snapshot copy to restore to on the destination after a break operation. This field is applicable only for SnapMirror relationships with FlexVol endpoints and when the PATCH state is "broken_off".
source	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.

Name	Type	Description
state	string	<p>State of the relationship. To initialize the relationship, PATCH the state to "snapmirrored". To break the relationship, PATCH the state to "broken_off". To resync the broken relationship, PATCH the state to "snapmirrored" for relationships with async policy type or "in_sync" for relationships with sync policy type. To pause the relationship, suspending further transfers, PATCH the state to "paused". To resume transfers for a paused relationship, PATCH the state to "snapmirrored" or "in_sync". The entries "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships of the sync policy type. 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" fields using GET to determine if the transition is successful.</p>
transfer	<a href="#">transfer</a>	Basic information on the current transfer.
unhealthy_reason	array[ <a href="#">snapmirror_error</a> ]	Reason the relationship is not healthy. It is a concatenation of up to four levels of error messages.
uuid	string	

## Example response

```
{
  "destination": {
    "path": "svm1:volumel",
    "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": "MirrorAndVault",
    "type": "async",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "source": {
    "path": "svm1:volumel",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "state": "broken_off",
  "transfer": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
}
```

```

    "state": "aborted",
    "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	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror\_endpoint

Endpoint of a SnapMirror relationship.

Name	Type	Description
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: <ul style="list-style-type: none"><li>• example: svm1:volume1</li><li>• readCreate: 1</li></ul>
svm	<a href="#">svm</a>	

policy

Basic policy information of the relationship.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
type	string	
uuid	string	

transfer

Basic information on the current transfer.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update a SnapMirror relationship

PATCH /snapmirror/relationships/{uuid}



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.

#### 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.

##### Performing a SnapMirror "resync"

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff/" '{"state":"snapmirrored"}'
```

##### Performing a SnapMirror "initialize"

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff/" '{"state":"snapmirrored"}'
```

##### Performing a SnapMirror "resume"

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff/" '{"state":"snapmirrored"}'
```

##### Performing a SnapMirror "quiesce"

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff" '{"state":"paused"}'
```

##### Performing a SnapMirror "break"

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff" '{"state":"broken_off"}'
```

## Updating an associated SnapMirror policy

```
PATCH "/api/snapmirror/relationships/9e922e65-1818-11e9-8b22-005056bbee73/" '{"policy": { "name" : "MirrorAndVaultDiscardNetwork"}}'
```

### Learn more

- [DOC /snapmirror/relationships](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Relationship UUID

### Request Body

Name	Type	Description
destination	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.
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	<a href="#">policy</a>	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 field is applicable only for relationships with FlexGroup or FlexVol endpoints and when the PATCH state is "snapmirrored".
quick_resync	boolean	Set to true to reduce resync time by not preserving storage efficiency. This field is applicable only for relationships with FlexVol endpoints and when the PATCH state is "snapmirrored".

Name	Type	Description
recover_after_break	boolean	Set to true to recover from a failed SnapMirror break operation on a FlexGroup relationship. This restores all destination FlexGroup constituents to the latest Snapshot copy, and any writes to the read-write constituents are lost. This field is applicable only for SnapMirror relationships with FlexGroup endpoints and when the PATCH state is "broken_off".
restore	boolean	Set to true to create a relationship for restore. To trigger restore-transfer, use transfers POST on the restore relationship.
restore_to_snapshot	string	Specifies the Snapshot copy to restore to on the destination after a break operation. This field is applicable only for SnapMirror relationships with FlexVol endpoints and when the PATCH state is "broken_off".
source	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.

Name	Type	Description
state	string	<p>State of the relationship. To initialize the relationship, PATCH the state to "snapmirrored". To break the relationship, PATCH the state to "broken_off". To resync the broken relationship, PATCH the state to "snapmirrored" for relationships with async policy type or "in_sync" for relationships with sync policy type. To pause the relationship, suspending further transfers, PATCH the state to "paused". To resume transfers for a paused relationship, PATCH the state to "snapmirrored" or "in_sync". The entries "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships of the sync policy type. 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" fields using GET to determine if the transition is successful.</p>
transfer	<a href="#">transfer</a>	Basic information on the current transfer.
unhealthy_reason	array[ <a href="#">snapmirror_error</a> ]	Reason the relationship is not healthy. It is a concatenation of up to four levels of error messages.
uuid	string	

## Example request

```
{
  "destination": {
    "path": "svm1:volumel",
    "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": "MirrorAndVault",
    "type": "async",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "source": {
    "path": "svm1:volumel",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "state": "broken_off",
  "transfer": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
}
```

```

    "state": "aborted",
    "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": {
    "_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
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

Error code	Description
13303811	Transition from paused state failed
13303820	SnapMirror policy was successfully updated, state transition failed

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	<a href="#">href</a>	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror\_endpoint

Endpoint of a SnapMirror relationship.

Name	Type	Description
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: <ul style="list-style-type: none"><li>• example: svm1:volume1</li><li>• readCreate: 1</li></ul>
svm	<a href="#">svm</a>	

policy

Basic policy information of the relationship.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
type	string	
uuid	string	

transfer

Basic information on the current transfer.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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

Name	Type	Description
destination	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.
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	<a href="#">policy</a>	Basic policy information of the relationship.

Name	Type	Description
preserve	boolean	Set to true on resync to preserve Snapshot copies on the destination that are newer than the latest common Snapshot copy. This field is applicable only for relationships with FlexGroup or FlexVol endpoints and when the PATCH state is "snapmirrored".
quick_resync	boolean	Set to true to reduce resync time by not preserving storage efficiency. This field is applicable only for relationships with FlexVol endpoints and when the PATCH state is "snapmirrored".
recover_after_break	boolean	Set to true to recover from a failed SnapMirror break operation on a FlexGroup relationship. This restores all destination FlexGroup constituents to the latest Snapshot copy, and any writes to the read-write constituents are lost. This field is applicable only for SnapMirror relationships with FlexGroup endpoints and when the PATCH state is "broken_off".
restore	boolean	Set to true to create a relationship for restore. To trigger restore-transfer, use transfers POST on the restore relationship.
restore_to_snapshot	string	Specifies the Snapshot copy to restore to on the destination after a break operation. This field is applicable only for SnapMirror relationships with FlexVol endpoints and when the PATCH state is "broken_off".
source	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.

Name	Type	Description
state	string	State of the relationship. To initialize the relationship, PATCH the state to "snapmirrored". To break the relationship, PATCH the state to "broken_off". To resync the broken relationship, PATCH the state to "snapmirrored" for relationships with async policy type or "in_sync" for relationships with sync policy type. To pause the relationship, suspending further transfers, PATCH the state to "paused". To resume transfers for a paused relationship, PATCH the state to "snapmirrored" or "in_sync". The entries "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships of the sync policy type. 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" fields using GET to determine if the transition is successful.
transfer	<a href="#">transfer</a>	Basic information on the current transfer.
unhealthy_reason	array[ <a href="#">snapmirror_error</a> ]	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
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## 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 and it 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

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
bytes_transferred	integer	query	False	Filter by bytes_transferred
snapshot	string	query	False	Filter by snapshot
uuid	string	query	False	Filter by uuid
checkpoint_size	integer	query	False	Filter by checkpoint_size
relationship.uuid	string	query	False	Filter by relationship.uuid
relationship.restore	boolean	query	False	Filter by relationship.restore
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
relationship.destination.path	string	query	False	Filter by relationship.destination.path
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">snapmirror_transfer</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "bytes_transferred": 0,
    "checkpoint_size": 0,
    "files": {
      "destination_path": "/dirb/file2",
      "source_path": "/dira/file1"
    },
    "relationship": {
      "destination": {
        "path": "svm1:volumel",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      },
      "name": "svm1",
      "uuid": "d2d7ceea-ab52-11e8-855e-00505682a4c7"
    },
    "snapshot": "string",
    "state": "aborted",
    "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
next	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror\_endpoint

Endpoint of a SnapMirror relationship.

Name	Type	Description
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: <ul style="list-style-type: none"> <li>• example: svm1:volume1</li> <li>• readCreate: 1</li> </ul>
svm	<a href="#">svm</a>	

#### relationship

Name	Type	Description
destination	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.
restore	boolean	Is the relationship for restore?
uuid	string	

#### snapmirror\_transfer

##### Snapmirror transfer information

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
bytes_transferred	integer	Bytes transferred
checkpoint_size	integer	Amount of data transferred in bytes as recorded in the restart checkpoint.
files	array[ <a href="#">files</a> ]	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	<a href="#">relationship</a>	
snapshot	string	Name of Snapshot copy being transferred.
source_snapshot	string	Specifies the Snapshot copy on the source to be transferred to the destination.

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.
storage_efficiency_enabled	boolean	This is supported for transfer of restore relationship only. Set this field 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[ <a href="#">error_arguments</a> ]	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

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.

### Performing a SnapMirror initialize or update

```
POST "/api/snapmirror/relationships/e4e7e130-0279-11e9-b566-0050568e9909/transfers" '{}'
```

### Performing a 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
return_timeout	integer	query	False	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.
relationship.uuid	string	path	True	Relationship UUID

#### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
bytes_transferred	integer	Bytes transferred
checkpoint_size	integer	Amount of data transferred in bytes as recorded in the restart checkpoint.
files	array[ <a href="#">files</a> ]	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	<a href="#">relationship</a>	

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.
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.
storage_efficiency_enabled	boolean	This is supported for transfer of restore relationship only. Set this field to 'false' to turn off storage efficiency for data transferred over the wire and written to the destination.
uuid	string	

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bytes_transferred": 0,
  "checkpoint_size": 0,
  "files": {
    "destination_path": "/dirb/file2",
    "source_path": "/dira/file1"
  },
  "relationship": {
    "destination": {
      "path": "svm1:volumel",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    },
    "name": "svm1",
    "uuid": "d2d7ceea-ab52-11e8-855e-00505682a4c7"
  },
  "snapshot": "string",
  "state": "aborted",
  "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

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	<a href="#">href</a>	

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	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror\_endpoint

Endpoint of a SnapMirror relationship.

Name	Type	Description
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: <ul style="list-style-type: none"><li>• example: svm1:volume1</li><li>• readCreate: 1</li></ul>
svm	<a href="#">svm</a>	

relationship

Name	Type	Description
destination	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.
restore	boolean	Is the relationship for restore?
uuid	string	

snapmirror\_transfer

Snapmirror transfer information

Name	Type	Description
_links	<a href="#">_links</a>	
bytes_transferred	integer	Bytes transferred
checkpoint_size	integer	Amount of data transferred in bytes as recorded in the restart checkpoint.
files	array[ <a href="#">files</a> ]	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	<a href="#">relationship</a>	
snapshot	string	Name of Snapshot copy being transferred.
source_snapshot	string	Specifies the Snapshot copy on the source to be transferred to the destination.
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.
storage_efficiency_enabled	boolean	This is supported for transfer of restore relationship only. Set this field to 'false' to turn off storage efficiency for data transferred over the wire and written to the destination.

Name	Type	Description
uuid	string	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}
```

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	<a href="#">_links</a>	
bytes_transferred	integer	Bytes transferred
checkpoint_size	integer	Amount of data transferred in bytes as recorded in the restart checkpoint.
files	array[ <a href="#">files</a> ]	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	<a href="#">relationship</a>	
snapshot	string	Name of Snapshot copy being transferred.
source_snapshot	string	Specifies the Snapshot copy on the source to be transferred to the destination.
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
storage_efficiency_enabled	boolean	This is supported for transfer of restore relationship only. Set this field to 'false' to turn off storage efficiency for data transferred over the wire and written to the destination.
uuid	string	

### Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bytes_transferred": 0,
  "checkpoint_size": 0,
  "files": {
    "destination_path": "/dirb/file2",
    "source_path": "/dira/file1"
  },
  "relationship": {
    "destination": {
      "path": "svm1:volumel",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "uuid": "d2d7ceea-ab52-11e8-855e-00505682a4c7"
},
"snapshot": "string",
"state": "aborted",
"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	<a href="#">href</a>	

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	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror\_endpoint

Endpoint of a SnapMirror relationship.

Name	Type	Description
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: <ul style="list-style-type: none"><li>• example: svm1:volume1</li><li>• readCreate: 1</li></ul>
svm	<a href="#">svm</a>	

relationship



Name	Type	Description
destination	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Cancel an ongoing SnapMirror transfer

PATCH /snapmirror/relationships/{relationship.uuid}/transfers/{uuid}

Aborts an ongoing SnapMirror transfer.

### 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
_links	<a href="#">_links</a>	
bytes_transferred	integer	Bytes transferred
checkpoint_size	integer	Amount of data transferred in bytes as recorded in the restart checkpoint.
files	array[ <a href="#">files</a> ]	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	<a href="#">relationship</a>	
snapshot	string	Name of Snapshot copy being transferred.
source_snapshot	string	Specifies the Snapshot copy on the source to be transferred to the destination.
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.
storage_efficiency_enabled	boolean	This is supported for transfer of restore relationship only. Set this field to 'false' to turn off storage efficiency for data transferred over the wire and written to the destination.

Name	Type	Description
uuid	string	

### Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bytes_transferred": 0,
  "checkpoint_size": 0,
  "files": {
    "destination_path": "/dirb/file2",
    "source_path": "/dira/file1"
  },
  "relationship": {
    "destination": {
      "path": "svm1:volumel",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "uuid": "d2d7ceea-ab52-11e8-855e-00505682a4c7"
},
"snapshot": "string",
"state": "aborted",
"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

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	<a href="#">href</a>	

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	

svm

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror\_endpoint

Endpoint of a SnapMirror relationship.

Name	Type	Description
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: <ul style="list-style-type: none"><li>• example: svm1:volume1</li><li>• readCreate: 1</li></ul>
svm	<a href="#">svm</a>	

relationship

Name	Type	Description
destination	<a href="#">snapmirror_endpoint</a>	Endpoint of a SnapMirror relationship.
restore	boolean	Is the relationship for restore?
uuid	string	

snapmirror\_transfer

Snapmirror transfer information

Name	Type	Description
_links	<a href="#">_links</a>	
bytes_transferred	integer	Bytes transferred
checkpoint_size	integer	Amount of data transferred in bytes as recorded in the restart checkpoint.
files	array[ <a href="#">files</a> ]	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	<a href="#">relationship</a>	
snapshot	string	Name of Snapshot copy being transferred.
source_snapshot	string	Specifies the Snapshot copy on the source to be transferred to the destination.
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.
storage_efficiency_enabled	boolean	This is supported for transfer of restore relationship only. Set this field to 'false' to turn off storage efficiency for data transferred over the wire and written to the destination.

Name	Type	Description
uuid	string	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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 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.

### 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.

### 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.

The list of patchable properties with a brief description for each is as follows:

- 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.

### **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 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 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-
b692-57d1ceec874d4" -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"
  },
  "hybrid_cache": {
```

```

    "enabled": false
  },
  "mirror": {
    "enabled": false,
    "state": "unmirrored"
  }
},
"plexes": [
  {
    "name": "plex0",
  }
],
"cloud_storage": {
  "attach_eligible": false
},
}

```

### 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"
      },
      "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"
      },
      "hybrid_cache": {
        "enabled": false
      },
      "mirror": {
        "enabled": false
      }
    },
  }
],
"num_records": 2,
}

```

## 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
    }
  },
}
```

### Retrieve a collection of aggregates for an entire cluster

GET /storage/aggregates

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 [DOC Requesting specific fields](#) to learn more.

- `space.block_storage.inactive_user_data`
- `space.footprint`

### Related ONTAP commands

- `storage aggregate show`

### Learn more

- [DOC /storage/aggregates](#)

### Parameters

Name	Type	In	Required	Description
<code>recommend</code>	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'.
<code>show_spare</code>	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'.
<code>data_encryption.software_encryption_enabled</code>	boolean	query	False	Filter by <code>data_encryption.software_encryption_enabled</code>
<code>data_encryption.drive_protection_enabled</code>	boolean	query	False	Filter by <code>data_encryption.drive_protection_enabled</code>



Name	Type	In	Required	Description
state	string	query	False	Filter by state
space.encyency_wit hout_snapshots.ratio	number	query	False	Filter by space.encyency_wit hout_snapshots.ratio
space.encyency_wit hout_snapshots.logical_used	integer	query	False	Filter by space.encyency_wit hout_snapshots.logical_used
space.encyency_wit hout_snapshots.savings	integer	query	False	Filter by space.encyency_wit hout_snapshots.savings
space.cloud_storage .used	integer	query	False	Filter by space.cloud_storage .used
space.footprint	integer	query	False	Filter by space.footprint
space.encyency.ratio	number	query	False	Filter by space.encyency.ratio
space.encyency.logical_used	integer	query	False	Filter by space.encyency.logical_used
space.encyency.savings	integer	query	False	Filter by space.encyency.savings
space.block_storage .size	integer	query	False	Filter by space.block_storage .size
space.block_storage .used	integer	query	False	Filter by space.block_storage .used
space.block_storage .full_threshold_percent	integer	query	False	Filter by space.block_storage .full_threshold_percent

Name	Type	In	Required	Description
space.block_storage.inactive_user_data	integer	query	False	Filter by space.block_storage.inactive_user_data
space.block_storage.available	integer	query	False	Filter by space.block_storage.available
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
name	string	query	False	Filter by name
home_node.name	string	query	False	Filter by home_node.name
home_node.uuid	string	query	False	Filter by home_node.uuid
uuid	string	query	False	Filter by uuid
snaplock_type	string	query	False	Filter by snaplock_type
block_storage.primary.checksum_style	string	query	False	Filter by block_storage.primary.checksum_style
block_storage.primary.raid_type	string	query	False	Filter by block_storage.primary.raid_type
block_storage.primary.raid_size	integer	query	False	Filter by block_storage.primary.raid_size
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

Name	Type	In	Required	Description
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.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.raid_type	string	query	False	Filter by block_storage.hybrid_cache.raid_type
block_storage.hybrid_cache.used	integer	query	False	Filter by block_storage.hybrid_cache.used
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
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
create_time	string	query	False	Filter by create_time
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
error	<a href="#">error</a>	
num_records	integer	Number of records
records	array[ <a href="#">aggregate</a> ]	
spares	array[ <a href="#">aggregate_spare</a> ]	

## 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": "raid_dp",
      "size": 1612709888,
      "used": 26501122
    },
    "mirror": {
      "enabled": "",
      "state": "unmirrored"
    },
    "plexes": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "plex0"
    },
    "primary": {
```

```

    "checksum_style": "block",
    "disk_class": "performance",
    "disk_count": 8,
    "raid_size": 16,
    "raid_type": "raid_dp"
  }
},
"cloud_storage": {
  "stores": {
    "cloud_store": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "store1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "used": 0
  }
},
"create_time": "2018-01-01 12:00:00 -0400",
"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"
},
"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"snaplock_type": "non_snaplock",

```

```

"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": "online",
"uuid": "string"
},
"spares": {
  "checksum_style": "block",
  "disk_class": "solid_state",
  "layout_requirements": {
    "aggregate_min_disks": 6,
    "raid_group": {
      "default": 16,
      "max": 28,
      "min": 5
    },
    "raid_type": "raid_dp"
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}

```

```

    },
    "size": 10156769280,
    "syncmirror_pool": "pool0",
    "usable": 9
  }
}

```

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
19726382	Another provisioning operation is in progress on this cluster. Wait a few minutes, and try the operation again.
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.
19726359	Aggregate recommendations are not supported on ONTAP Select.
19726341	Not enough eligible spare disks are available on the node.
787092	The target field cannot be specified for this 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

\_links

Name	Type	Description
self	<a href="#">href</a>	

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	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
<a href="#">_links</a>	<a href="#">_links</a>	
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.
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	<a href="#">hybrid_cache</a>	Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.
mirror	<a href="#">mirror</a>	
plexes	array[ <a href="#">plex_reference</a> ]	Plex reference for each plex in the aggregate.
primary	<a href="#">primary</a>	Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

### cloud\_store

Cloud store

Name	Type	Description
_links	<a href="#">_links</a>	
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	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	Aggregate uses self-encrypting drives with data protection enabled.
software_encryption_enabled	boolean	NetApp Aggregate Encryption enabled. All data in the aggregate is encrypted.

#### dr\_home\_node

Node where the aggregate belongs 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 belongs after giveback. The value for this field might differ from the value of the 'node' field during takeover.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

## node

Node where the aggregate currently resides.

Name	Type	Description
_links	<a href="#">_links</a>	
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> .

Name	Type	Description
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 aggregate 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	<a href="#">block_storage</a>	
cloud_storage	<a href="#">cloud_storage</a>	
efficiency	<a href="#">efficiency</a>	Storage efficiency
efficiency_without_snapshots	<a href="#">efficiency_without_snapshots</a>	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 <code>block_storage</code> tier and the <code>cloud_storage</code> 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 <code>footprint</code> or <code>**</code> .

#### aggregate

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
block_storage	<a href="#">block_storage</a>	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	<a href="#">cloud_storage</a>	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	<a href="#">data_encryption</a>	



Name	Type	Description
dr_home_node	<a href="#">dr_home_node</a>	Node where the aggregate belongs after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	<a href="#">home_node</a>	Node where the aggregate belongs after giveback. The value for this field might differ from the value of the 'node' field during takeover.
name	string	Aggregate name
node	<a href="#">node</a>	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	<a href="#">space</a>	
state	string	Operational state of the aggregate
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

Name	Type	Description
default	boolean	Indicates if this RAID type is the default
raid_group	<a href="#">raid_group</a>	
raid_type	string	RAID type

node

Node where the spares are assigned

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

aggregate\_spare

Name	Type	Description
checksum_style	string	The checksum type that has been assigned to the spares
disk_class	string	Disk class of spares
layout_requirements	array[ <a href="#">layout_requirement</a> ]	Available RAID protections and their restrictions
node	<a href="#">node</a>	Node where the spares are assigned
size	integer	Usable size of each spare in bytes
syncmirror_pool	string	SyncMirror spare pool
usable	integer	Total number of usable spares

## Create a collection of aggregates for an entire cluster

POST /storage/aggregates

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.

## 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"}}
```

## Learn more

- [DOC /storage/aggregates](#)

## Parameters

Name	Type	In	Required	Description
<code>disk_size</code>	integer	query	False	If set, POST only selects disks of the specified size.

## Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	

Name	Type	Description
block_storage	<a href="#">block_storage</a>	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	<a href="#">cloud_storage</a>	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	<a href="#">data_encryption</a>	
dr_home_node	<a href="#">dr_home_node</a>	Node where the aggregate belongs after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	<a href="#">home_node</a>	Node where the aggregate belongs after giveback. The value for this field might differ from the value of the 'node' field during takeover.
name	string	Aggregate name
node	<a href="#">node</a>	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	<a href="#">space</a>	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "raid_type": "raid_dp",
      "size": 1612709888,
      "used": 26501122
    },
    "mirror": {
      "enabled": "",
      "state": "unmirrored"
    },
    "plexes": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "plex0"
    },
    "primary": {
      "checksum_style": "block",
      "disk_class": "performance",
      "disk_count": 8,
      "raid_size": 16,
      "raid_type": "raid_dp"
    }
  },
  "cloud_storage": {
    "stores": {
      "cloud_store": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "store1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  },
}
```

```

    "used": 0
  }
},
"create_time": "2018-01-01 12:00:00 -0400",
"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"
},
"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"snaplock_type": "non_snaplock",
"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": "online",
"uuid": "string"
}

```

### 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
787092	The target field cannot be specified for this operation.
786468	VLDB is offline.
786439	An aggregate already uses the specified name.

Error Code	Description
786819	The value is invalid for the specified option at the current privilege level.
786438	Failed to create an aggregate on the node.
786902	RAID-TEC aggregate is not fully supported.
786911	Not every node in the cluster has the Data ONTAP version required for the feature.
786446	The node is not in cluster.
787069	Node is setup for MetroCluster over IP configuration; creating an unmirrored aggregate is not supported in this configuration.
19726344	No recommendation can be made for this cluster.
19726373	Recommendation specified for creating aggregates is not current.
2425736	No matching node found for the target UUID.
1114292	The required SnapLock license is not installed.
19726382	Another provisioning operation is in progress on this cluster. Wait a few minutes, and try the operation again.
19726357	Automatic aggregate creation is not supported on MetroCluster.
19726358	Automatic aggregate creation is not supported on ONTAP Cloud.
19726359	Automatic aggregate creation is not supported on ONTAP Select.
19726341	Not enough eligible spare disks are available on the node.
460770	The aggregate create job failed to create the 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	<a href="#">href</a>	

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 <code>hybrid_cache.enabled</code> is 'true'.
enabled	boolean	Aggregate uses HDDs with SSDs as a cache
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when <code>hybrid_cache.enabled</code> is 'true'.
size	integer	Total usable space in bytes of SSD cache. Only provided when <code>hybrid_cache.enabled</code> is 'true'.
used	integer	Space used in bytes of SSD cache. Only provided when <code>hybrid_cache.enabled</code> is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex\_reference

## Plex

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>name</code>	string	

### primary

Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

Name	Type	Description
<code>checksum_style</code>	string	The checksum style used by the aggregate.
<code>disk_class</code>	string	The class of disks being used by the aggregate.
<code>disk_count</code>	integer	Number of disks used in the aggregate. This includes parity disks, but excludes disks in the hybrid cache.
<code>raid_size</code>	integer	Option to specify the maximum number of disks that can be included in a RAID group.
<code>raid_type</code>	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
<code>hybrid_cache</code>	<a href="#">hybrid_cache</a>	Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.
<code>mirror</code>	<a href="#">mirror</a>	
<code>plexes</code>	array[ <a href="#">plex_reference</a> ]	Plex reference for each plex in the aggregate.

Name	Type	Description
primary	<a href="#">primary</a>	Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

cloud\_store

Cloud store

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

cloud\_storage\_tier

Name	Type	Description
cloud_store	<a href="#">cloud_store</a>	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	Aggregate is eligible for a cloud store to be attached.
stores	array[ <a href="#">cloud_storage_tier</a> ]	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	Aggregate uses self-encrypting drives with data protection enabled.
software_encryption_enabled	boolean	NetApp Aggregate Encryption enabled. All data in the aggregate is encrypted.

dr\_home\_node

Node where the aggregate belongs 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 belongs after giveback. The value for this field might differ from the value of the 'node' field during takeover.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

node

Node where the aggregate currently resides.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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 <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 aggregate 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	<a href="#">block_storage</a>	
cloud_storage	<a href="#">cloud_storage</a>	
efficiency	<a href="#">efficiency</a>	Storage efficiency
efficiency_without_snapshots	<a href="#">efficiency_without_snapshots</a>	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 **.

## aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
block_storage	<a href="#">block_storage</a>	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	<a href="#">cloud_storage</a>	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	<a href="#">data_encryption</a>	
dr_home_node	<a href="#">dr_home_node</a>	Node where the aggregate belongs after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	<a href="#">home_node</a>	Node where the aggregate belongs after giveback. The value for this field might differ from the value of the 'node' field during takeover.
name	string	Aggregate name
node	<a href="#">node</a>	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	<a href="#">space</a>	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

job\_link



Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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

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
unreclaimed_space_threshold	integer	query	False	Filter by unreclaimed_space_threshold

Name	Type	In	Required	Description
mirror_degraded	boolean	query	False	Filter by mirror_degraded
target.uuid	string	query	False	Filter by target.uuid
target.name	string	query	False	Filter by target.name
used	integer	query	False	Filter by used
primary	boolean	query	False	Filter by primary
availability	string	query	False	Filter by availability
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">cloud_store</a> ]	

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "availability": "available",
    "target": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "target1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "unreclaimed_space_threshold": 20,
    "used": 0
  }
}
```

## Error

Status: Default, n/a

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

target

Cloud target

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

cloud\_store

Name	Type	Description
_links	<a href="#">_links</a>	
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 cloud mirrored composite aggregate.
target	<a href="#">target</a>	Cloud target

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.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

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
<code>aggregate.uuid</code>	string	path	True	Aggregate UUID
<code>allow_flexgroups</code>	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.
<code>check_only</code>	boolean	query	False	Validate only and do not add the cloud store.

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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 cloud mirrored composite aggregate.
target	<a href="#">target</a>	Cloud target
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

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "availability": "available",
  "target": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "target1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "unreclaimed_space_threshold": 20,
  "used": 0
}
```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

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

## Error

Status: Default, n/a

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	<a href="#">href</a>	

target

Cloud target

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

cloud\_store

Name	Type	Description
_links	<a href="#">_links</a>	
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 cloud mirrored composite aggregate.
target	<a href="#">target</a>	Cloud target
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
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}
```

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	
target.uuid	string	path	True	

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

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

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">cloud_store</a> ]	



## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "availability": "available",
    "target": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "target1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "unreclaimed_space_threshold": 20,
    "used": 0
  }
}
```

## Error

Status: Default, n/a

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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

target

Cloud target

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

cloud\_store

Name	Type	Description
_links	<a href="#">_links</a>	
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 cloud mirrored composite aggregate.
target	<a href="#">target</a>	Cloud target

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.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update a cloud store

PATCH /storage/aggregates/{aggregate.uuid}/cloud-stores/{target.uuid}

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

### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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 cloud mirrored composite aggregate.
target	<a href="#">target</a>	Cloud target
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

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "availability": "available",
  "target": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "target1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "unreclaimed_space_threshold": 20,
  "used": 0
}
```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

target

Cloud target

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

cloud\_store

Name	Type	Description
_links	<a href="#">_links</a>	
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 cloud mirrored composite aggregate.
target	<a href="#">target</a>	Cloud target
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
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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

Retrieves the collection of plexes for the specified aggregate.

### Related ONTAP commands

- `storage aggregate plex show`

### Learn more

- [DOC /storage/aggregates/{aggregate.uuid}/plexes](#)

### Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID

Name	Type	In	Required	Description
aggregate.uuid	string	query	False	Filter by aggregate.uuid
aggregate.name	string	query	False	Filter by aggregate.name
name	string	path	False	Filter by name
resync.percent	integer	query	False	Filter by resync.percent
resync.level	string	query	False	Filter by resync.level
resync.active	boolean	query	False	Filter by resync.active
online	boolean	query	False	Filter by online
state	string	query	False	Filter by state
pool	string	query	False	Filter by pool
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.
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.

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	<a href="#">_links</a>	
error	<a href="#">error</a>	
num_records	integer	Number of records
records	array[ <a href="#">plex</a> ]	

## 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-ae1c-123478563412"
    },
    "name": "plex0",
    "pool": "pool0",
    "raid_groups": {
      "disks": {
        "disk": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "1.0.1"
        },
        "position": "data",
        "state": "normal",
        "type": "ssd",
        "usable_size": 947912704
      }
    }
  }
}
```

```

    },
    "name": "rg0",
    "recomputing_parity": {
      "percent": 10
    },
    "reconstruct": {
      "percent": 10
    }
  },
  "resync": {
    "level": "full",
    "percent": 10
  },
  "state": "normal"
}

```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

\_links

Name	Type	Description
self	<a href="#">href</a>	

aggregate

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	

Name	Type	Description
uuid	string	

disk

Disk

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	

raid\_group\_disk

Name	Type	Description
disk	<a href="#">disk</a>	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}

Retrieves the plex specified by the aggregate UUID and plex name.

#### Related ONTAP commands

- `storage aggregate plex show`

#### Learn more

- [DOC /storage/aggregates/{aggregate.uuid}/plexes](#)

#### 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	<a href="#">aggregate</a>	Aggregate
name	string	Plex name
online	boolean	Plex is online
pool	string	SyncMirror pool assignment
raid_groups	array[ <a href="#">raid_group</a> ]	
resync	<a href="#">resync</a>	
state	string	Plex state

## Example response

```
{
  "aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "plex0",
  "pool": "pool0",
  "raid_groups": {
    "disks": {
      "disk": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "1.0.1"
      },
      "position": "data",
      "state": "normal",
      "type": "ssd",
      "usable_size": 947912704
    },
    "name": "rg0",
    "recomputing_parity": {
      "percent": 10
    },
    "reconstruct": {
      "percent": 10
    }
  },
  "resync": {
    "level": "full",
    "percent": 10
  },
  "state": "normal"
}
```

## 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	<a href="#">href</a>	

aggregate

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

disk

Disk

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	

raid\_group\_disk

Name	Type	Description
disk	<a href="#">disk</a>	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	<a href="#">recomputing_parity</a>	
reconstruct	<a href="#">reconstruct</a>	

### 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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Delete an aggregate specified by the UUID

DELETE /storage/aggregates/{uuid}

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`

### Learn more

- [DOC /storage/aggregates](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

### Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

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

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
786468	VLDB is offline.
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.
786497	Cannot delete an aggregate that has volumes.
786435	Internal Error. Failed to create a communication handle.
786451	Failed to delete specified aggregate.
786472	Node that hosts the aggregate is offline.
460777	Failed to get information on the delete job.
460770	The aggregate delete job failed to delete the 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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 [DOC Requesting specific fields](#) to learn more.

- `space.block_storage.inactive_user_data`
- `space.footprint`

### Related ONTAP commands

- `storage aggregate show`

### Learn more

- [DOC /storage/aggregates](#)

### Parameters

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

### Response

```
Status: 200, Ok
```

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>block_storage</code>	<a href="#">block_storage</a>	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.
<code>cloud_storage</code>	<a href="#">cloud_storage</a>	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
<code>create_time</code>	string	Timestamp of aggregate creation

Name	Type	Description
data_encryption	<a href="#">data_encryption</a>	
dr_home_node	<a href="#">dr_home_node</a>	Node where the aggregate belongs after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	<a href="#">home_node</a>	Node where the aggregate belongs after giveback. The value for this field might differ from the value of the 'node' field during takeover.
name	string	Aggregate name
node	<a href="#">node</a>	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	<a href="#">space</a>	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "raid_type": "raid_dp",
      "size": 1612709888,
      "used": 26501122
    },
    "mirror": {
      "enabled": "",
      "state": "unmirrored"
    },
    "plexes": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "plex0"
    },
    "primary": {
      "checksum_style": "block",
      "disk_class": "performance",
      "disk_count": 8,
      "raid_size": 16,
      "raid_type": "raid_dp"
    }
  },
  "cloud_storage": {
    "stores": {
      "cloud_store": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "store1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  },
}
```

```

    "used": 0
  }
},
"create_time": "2018-01-01 12:00:00 -0400",
"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"
},
"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"snaplock_type": "non_snaplock",
"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": "online",
"uuid": "string"
}
```

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
787092	The target field cannot be specified for this 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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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	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
<code>_links</code>	<a href="#">_links</a>	
<code>name</code>	string	

### primary

Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

Name	Type	Description
<code>checksum_style</code>	string	The checksum style used by the aggregate.
<code>disk_class</code>	string	The class of disks being used by the aggregate.
<code>disk_count</code>	integer	Number of disks used in the aggregate. This includes parity disks, but excludes disks in the hybrid cache.
<code>raid_size</code>	integer	Option to specify the maximum number of disks that can be included in a RAID group.
<code>raid_type</code>	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
<code>hybrid_cache</code>	<a href="#">hybrid_cache</a>	Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.
<code>mirror</code>	<a href="#">mirror</a>	
<code>plexes</code>	array[ <a href="#">plex_reference</a> ]	Plex reference for each plex in the aggregate.

Name	Type	Description
primary	<a href="#">primary</a>	Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

cloud\_store

Cloud store

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

cloud\_storage\_tier

Name	Type	Description
cloud_store	<a href="#">cloud_store</a>	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	Aggregate is eligible for a cloud store to be attached.
stores	array[ <a href="#">cloud_storage_tier</a> ]	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	Aggregate uses self-encrypting drives with data protection enabled.
software_encryption_enabled	boolean	NetApp Aggregate Encryption enabled. All data in the aggregate is encrypted.

dr\_home\_node

Node where the aggregate belongs 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 belongs after giveback. The value for this field might differ from the value of the 'node' field during takeover.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

node

Node where the aggregate currently resides.

Name	Type	Description
_links	<a href="#">_links</a>	
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 <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 aggregate 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	<a href="#">block_storage</a>	
cloud_storage	<a href="#">cloud_storage</a>	
efficiency	<a href="#">efficiency</a>	Storage efficiency
efficiency_without_snapshots	<a href="#">efficiency_without_snapshots</a>	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 **.

## error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

## error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
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}

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`

### Learn more

- [DOC /storage/aggregates](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Aggregate UUID



Name	Type	In	Required	Description
disk_size	integer	query	False	If set, PATCH only selects disks of the specified size.

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
block_storage	<a href="#">block_storage</a>	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	<a href="#">cloud_storage</a>	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	<a href="#">data_encryption</a>	
dr_home_node	<a href="#">dr_home_node</a>	Node where the aggregate belongs after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	<a href="#">home_node</a>	Node where the aggregate belongs after giveback. The value for this field might differ from the value of the 'node' field during takeover.
name	string	Aggregate name
node	<a href="#">node</a>	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	<a href="#">space</a>	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "raid_type": "raid_dp",
      "size": 1612709888,
      "used": 26501122
    },
    "mirror": {
      "enabled": "",
      "state": "unmirrored"
    },
    "plexes": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "plex0"
    },
    "primary": {
      "checksum_style": "block",
      "disk_class": "performance",
      "disk_count": 8,
      "raid_size": 16,
      "raid_type": "raid_dp"
    }
  },
  "cloud_storage": {
    "stores": {
      "cloud_store": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "store1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  },
}
```

```

    "used": 0
  }
},
"create_time": "2018-01-01 12:00:00 -0400",
"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"
},
"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"snaplock_type": "non_snaplock",
"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": "online",
"uuid": "string"
}

```

## 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
787092	The target field cannot be specified for this operation.
787169	Only one field can be modified per operation.
786439	An aggregate already uses the specified name.

Error Code	Description
786468	VLDB is offline.
786771	Aggregate does not exist.
786458	Failed to rename aggregate.
786435	Internal Error. Failed to create a communication handle.
786867	Specified aggregate resides on the remote cluster.
786472	Node that hosts the aggregate is offline.
786434	Cannot connect to node where the aggregate resides.
786730	Internal Error
786911	Not every node in the cluster has the Data ONTAP version required for the feature.
787170	Failed to patch the "block_storage.primary.disk_count" because the disk count specified is smaller than existing disk count.
786965	Spare Selection in userspace failed.
786456	Failed to add disks to aggregate.
787046	Mirroring of a FabricPool is not allowed.
786808	Aggregate mirror failed.
786479	Cannot find node ID for the node.
787144	Aggregate is not a FabricPool.
786787	Aggregate is not online.
787156	Modifying the attributes of mirror object store is not allowed.
786923	This operation is disallowed during pre-commit phase of 7-mode to clustered Data ONTAP transition.
786924	Internal Error for an aggregate that is in 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.
786447	Failed to modify aggregate.
787178	Unmirroring an aggregate with a PATCH operation is not supported.
460777	Failed to get information on the job.
2425736	No matching node found for the UUID provided.

Error Code	Description
26542101	Unable to contact source node.
26542102	Unable to contact destination node.
26542084	Source node is at higher Data ONTAP version than destination node.
26542083	Destination node is at higher Data ONTAP version than source node.
26542120	A Vserver migrate operation is in progress. When the migrate operation completes, try the operation again.
26542121	A MetroCluster disaster recovery operation is in progress. When the recovery operation completes, try the operation again.
26542097	Unable to get D-blade ID of destination.
262247	The value is invalid for the field.
13108106	Cannot run aggregate relocation because volume expand is in progress.

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	<a href="#">href</a>	

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	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
<code>_links</code>	<a href="#">_links</a>	
<code>name</code>	string	

### primary

Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

Name	Type	Description
<code>checksum_style</code>	string	The checksum style used by the aggregate.
<code>disk_class</code>	string	The class of disks being used by the aggregate.
<code>disk_count</code>	integer	Number of disks used in the aggregate. This includes parity disks, but excludes disks in the hybrid cache.
<code>raid_size</code>	integer	Option to specify the maximum number of disks that can be included in a RAID group.
<code>raid_type</code>	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
<code>hybrid_cache</code>	<a href="#">hybrid_cache</a>	Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.
<code>mirror</code>	<a href="#">mirror</a>	
<code>plexes</code>	array[ <a href="#">plex_reference</a> ]	Plex reference for each plex in the aggregate.



Name	Type	Description
primary	<a href="#">primary</a>	Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

cloud\_store

Cloud store

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	
uuid	string	

cloud\_storage\_tier

Name	Type	Description
cloud_store	<a href="#">cloud_store</a>	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	Aggregate is eligible for a cloud store to be attached.
stores	array[ <a href="#">cloud_storage_tier</a> ]	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	Aggregate uses self-encrypting drives with data protection enabled.
software_encryption_enabled	boolean	NetApp Aggregate Encryption enabled. All data in the aggregate is encrypted.

dr\_home\_node

Node where the aggregate belongs 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 belongs after giveback. The value for this field might differ from the value of the 'node' field during takeover.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

node

Node where the aggregate currently resides.

Name	Type	Description
_links	<a href="#">_links</a>	
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 <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 aggregate 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	<a href="#">block_storage</a>	
cloud_storage	<a href="#">cloud_storage</a>	
efficiency	<a href="#">efficiency</a>	Storage efficiency
efficiency_without_snapshots	<a href="#">efficiency_without_snapshots</a>	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 **.

aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
block_storage	<a href="#">block_storage</a>	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	<a href="#">cloud_storage</a>	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	<a href="#">data_encryption</a>	
dr_home_node	<a href="#">dr_home_node</a>	Node where the aggregate belongs after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	<a href="#">home_node</a>	Node where the aggregate belongs after giveback. The value for this field might differ from the value of the 'node' field during takeover.
name	string	Aggregate name
node	<a href="#">node</a>	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	<a href="#">space</a>	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

job\_link

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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

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
fields	array[string]	query	False	Specify the fields to return.

### 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": {
      "type": "hdd"
    }
  },
  "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

### medias

Name	Type	Description
available	integer	Available space
size	integer	Total space
type	string	The type of media being used
used	integer	Used space

### block\_storage

Name	Type	Description
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.

### 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)

### error\_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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":
  "50000394:0808AA88:00000000:00000000:00000000:00000000:00000000:00000000:0
  0000000:00000000",
  "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"
    }
}
}
}

```

## Retrieve a collection of disks

GET /storage/disks

Retrieves a collection of disks.

#### Related ONTAP commands

- `storage disk show`

#### Learn more

- [DOC /storage/disks](#)

#### Parameters

Name	Type	In	Required	Description
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
bay	string	query	False	Filter by bay
vendor	string	query	False	Filter by vendor
dr_node.uuid	string	query	False	Filter by dr_node.uuid
dr_node.name	string	query	False	Filter by dr_node.name
rpm	integer	query	False	Filter by rpm
serial_number	string	query	False	Filter by serial_number
shelf.uuid	string	query	False	Filter by shelf.uuid
pool	string	query	False	Filter by pool
usable_size	integer	query	False	Filter by usable_size
home_node.name	string	query	False	Filter by home_node.name
home_node.uuid	string	query	False	Filter by home_node.uuid
model	string	query	False	Filter by model

Name	Type	In	Required	Description
aggregates.uuid	string	query	False	Filter by aggregates.uuid
aggregates.name	string	query	False	Filter by aggregates.name
type	string	query	False	Filter by type
state	string	query	False	Filter by state
name	string	query	False	Filter by name
rated_life_used_percent	integer	query	False	Filter by rated_life_used_percent
firmware_version	string	query	False	Filter by firmware_version
uid	string	query	False	Filter by uid
class	string	query	False	Filter by class
drawer.slot	integer	query	False	Filter by drawer.slot
drawer.id	integer	query	False	Filter by drawer.id
container_type	string	query	False	Filter by container_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.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">disk</a> ]	



## 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"
    },
    "firmware_version": "NA51",
    "home_node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "model": "X421_HCOBE450A10",
    "name": "1.0.1",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  },
}
```

```

    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "pool": "pool0",
  "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	

Name	Type	Description
name	string	
uuid	string	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

shelf\_reference

Shelf

Name	Type	Description
_links	<a href="#">_links</a>	
uid	string	

disk

Name	Type	Description
aggregates	array[ <a href="#">aggregates</a> ]	List of aggregates sharing this disk
bay	string	Disk shelf bay
class	string	Disk class
container_type	string	Type of overlying disk container
dr_node	<a href="#">dr_node</a>	
drawer	<a href="#">drawer</a>	
firmware_version	string	
home_node	<a href="#">home_node</a>	
model	string	
name	string	Cluster-wide disk name
node	<a href="#">node</a>	
pool	string	Pool to which disk is assigned

Name	Type	Description
rated_life_used_percent	integer	Percentage of rated life used
rpm	integer	Revolutions per minute
serial_number	string	
shelf	<a href="#">shelf_reference</a>	Shelf
state	string	State
type	string	Disk interface type
uid	string	The unique identifier for a disk
usable_size	integer	
vendor	string	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Retrieve a specific disk

GET /storage/disks/{name}

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[ <a href="#">aggregates</a> ]	List of aggregates sharing this disk
bay	string	Disk shelf bay
class	string	Disk class
container_type	string	Type of overlying disk container
dr_node	<a href="#">dr_node</a>	
drawer	<a href="#">drawer</a>	
firmware_version	string	
home_node	<a href="#">home_node</a>	
model	string	
name	string	Cluster-wide disk name
node	<a href="#">node</a>	
pool	string	Pool to which disk is assigned
rated_life_used_percent	integer	Percentage of rated life used
rpm	integer	Revolutions per minute

Name	Type	Description
serial_number	string	
shelf	<a href="#">shelf_reference</a>	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"
  },
  "firmware_version": "NA51",
  "home_node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "model": "X421_HCOBE450A10",
  "name": "1.0.1",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "pool": "pool0",
  "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
0:00000000:00000000",
  "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	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
name	string	
uuid	string	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

shelf\_reference

Shelf

Name	Type	Description
_links	<a href="#">_links</a>	
uid	string	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Create a clone of the file

POST /storage/file/clone

Creates a clone of the file.

### Request Body

Name	Type	Description
autodelete	boolean	Mark clone file for auto deletion.

Name	Type	Description
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	<a href="#">volume</a>	

### Example request

```
{
  "destination_path": "dest_file1, dir1/dest_file2",
  "range": [
    36605,
    73210
  ],
  "source_path": "src_file1, dir1/src_file2,
  ../.snapshot/snap1/src_file3",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

### 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	<a href="#">href</a>	

volume

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li></ul>

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"

Name	Type	Description
source_path	string	Relative path of the source file in the volume.
volume	<a href="#">volume</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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:



&ndash; Intra-Vserver where FlexCache and the corresponding origin volume reside in the same Vserver.

&ndash; Cross-Vserver but intra-cluster where FlexCache and the origin volume reside in the same cluster but belong to different Vservers.

&ndash; 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:

&ndash; GET /api/storage/flexcache/flexcaches

&ndash; GET /api/storage/flexcache/flexcaches/{uuid}

&ndash; POST /api/storage/flexcache/flexcaches

&ndash; 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`, `aggregates`, `path`, `origins.ip_address`, `origins.size`, `origins.state` are expensive to retrieve. The recommended method to use this API is to use 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, aggregates, path, origins.ip\_address, origins.size, origins.state are expensive to retrieve. The recommended method to use this API is to use 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,
"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_01",
      "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

Retrieves FlexCaches 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 [DOC 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.
- `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
aggregates.uuid	string	query	False	Filter by aggregates.uuid
aggregates.name	string	query	False	Filter by aggregates.name
uuid	string	query	False	Filter by uuid
path	string	query	False	Filter by path
size	integer	query	False	Filter by size
constituents_per_aggregate	integer	query	False	Filter by constituents_per_aggregate
origins.cluster.uuid	string	query	False	Filter by origins.cluster.uuid
origins.cluster.name	string	query	False	Filter by origins.cluster.name
origins.ip_address	string	query	False	Filter by origins.ip_address
origins.state	string	query	False	Filter by origins.state
origins.volume.name	string	query	False	Filter by origins.volume.name
origins.volume.uuid	string	query	False	Filter by origins.volume.uuid
origins.size	integer	query	False	Filter by origins.size
origins.svm.uuid	string	query	False	Filter by origins.svm.uuid
origins.svm.name	string	query	False	Filter by origins.svm.name
origins.create_time	string	query	False	Filter by origins.create_time

Name	Type	In	Required	Description
name	string	query	False	Filter by name
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
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.
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	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">flexcache</a> ]	

## 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"
    },
    "name": "voll",
    "origins": {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "create_time": "2018-06-04 19:00:00 UTC",
      "ip_address": "10.10.10.7",
      "size": 0,
      "state": "error",
      "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

cluster

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

## volume

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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"><li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li></ul>

## flexcache\_relationship

Name	Type	Description
<a href="#">cluster</a>	<a href="#">cluster</a>	
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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
<a href="#">volume</a>	<a href="#">volume</a>	

## svm

### FlexCache SVM

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
name	string	FlexCache name
origins	array[ <a href="#">flexcache_relationship</a> ]	
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	<a href="#">svm</a>	FlexCache SVM
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[ <a href="#">error_arguments</a> ]	Message arguments

Name	Type	Description
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

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.
- constituents\_per\_aggregate - 4 if aggregates.name or aggregates.uuid is used.

#### Related ONTAP commands

- volume flexcache create

#### Learn more

- [DOC /storage/flexcache/flexcaches](#)

#### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	

Name	Type	Description
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
name	string	FlexCache name
origins	array[ <a href="#">flexcache_relationship</a> ]	
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	<a href="#">svm</a>	FlexCache SVM
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "voll",
  "origins": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "create_time": "2018-06-04 19:00:00 UTC",
    "ip_address": "10.10.10.7",
    "size": 0,
    "state": "error",
    "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"
}

```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

```

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

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

cluster

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	<a href="#">_links</a>	

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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

#### flexcache\_relationship

Name	Type	Description
cluster	<a href="#">cluster</a>	
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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
volume	<a href="#">volume</a>	

#### svm

##### FlexCache SVM

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
name	string	FlexCache name
origins	array[ <a href="#">flexcache_relationship</a> ]	
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	<a href="#">svm</a>	FlexCache SVM
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.

#### job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}
```

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	

#### Response

```
Status: 202, Accepted
```

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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 use filter and retrieve only the required fields. See [DOC 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.
- `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
<code>uuid</code>	string	path	True	Unique identifier of FlexCache.
<code>fields</code>	array[string]	query	False	Specify the fields to return.

### Response

```
Status: 200, Ok
```

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>aggregates</code>	array[ <a href="#">aggregates</a> ]	
<code>constituents_per_aggregate</code>	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
<code>name</code>	string	FlexCache name

Name	Type	Description
origins	array[ <a href="#">flexcache_relationship</a> ]	
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	<a href="#">svm</a>	FlexCache SVM
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"
  },
  "name": "voll",
  "origins": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "create_time": "2018-06-04 19:00:00 UTC",
    "ip_address": "10.10.10.7",
    "size": 0,
    "state": "error",
    "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	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

cluster

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	<a href="#">_links</a>	

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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

#### flexcache\_relationship

Name	Type	Description
cluster	<a href="#">cluster</a>	
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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
volume	<a href="#">volume</a>	

#### svm

##### FlexCache SVM

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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:

- &ndash; Intra-Vserver where FlexCache and the corresponding origin volume reside in the same Vserver.
- &ndash; Cross-Vserver but intra-cluster where FlexCache and the origin volume reside in the same cluster but belong to different Vservers.
- &ndash; 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:

- &ndash; GET `/api/storage/flexcache/origins`
- &ndash; 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/80fcae4-
0dc2-488b-afb8-86d28a34cda8" -H "accept: application/json"

# The response:
{
  "uuid": "80fcae4-0dc2-488b-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-aed0-0050568eddb"
      },
      "svm": {
        "name": "vs_1_4",
        "uuid": "36f68322-0e93-11e9-aed0-0050568eddb"
      },
      "cluster": {
        "name": "node4",
        "uuid": "c32f16b8-0e90-11e9-aed0-0050568eddb"
      }
    },
    {
      "ip_address": "10.140.103.183",
      "create_time": "2019-01-02T21:08:34+05:30",
      "volume": {
        "name": "fc_421",
        "uuid": "71ee8f36-0ea4-11e9-aed0-0050568eddb"
      },
      "svm": {
        "name": "vs_1_4",
        "uuid": "36f68322-0e93-11e9-aed0-0050568eddb"
      },
      "cluster": {
        "name": "node4",

```

```

    "uuid": "c32f16b8-0e90-11e9-aed0-0050568eddbb"
  }
},
{
  "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-aed0-0050568eddbb"
  },
  "cluster": {
    "name": "node4",
    "uuid": "c32f16b8-0e90-11e9-aed0-0050568eddbb"
  }
},
{
  "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

Retrieves origin of FlexCaches 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 [DOC Requesting specific fields](#) to learn more.

- `flexcaches.ip_address` - IP address of FlexCache.
- `flexcaches.size` - Physical size 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
flexcaches.cluster.uuid	string	query	False	Filter by flexcaches.cluster.uuid
flexcaches.cluster.name	string	query	False	Filter by flexcaches.cluster.name
flexcaches.ip_addresses	string	query	False	Filter by flexcaches.ip_addresses

Name	Type	In	Required	Description
flexcaches.state	string	query	False	Filter by flexcaches.state
flexcaches.volume.name	string	query	False	Filter by flexcaches.volume.name
flexcaches.volume.uuid	string	query	False	Filter by flexcaches.volume.uuid
flexcaches.size	integer	query	False	Filter by flexcaches.size
flexcaches.svm.uuid	string	query	False	Filter by flexcaches.svm.uuid
flexcaches.svm.name	string	query	False	Filter by flexcaches.svm.name
flexcaches.create_time	string	query	False	Filter by flexcaches.create_time
name	string	query	False	Filter by name
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
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.
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.
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	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">flexcache_origin</a> ]	



## 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": "2018-06-04 19:00:00 UTC",
      "ip_address": "10.10.10.7",
      "size": 0,
      "state": "error",
      "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": "voll, 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

cluster

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the volume.

Name	Type	Description
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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

#### flexcache\_relationship

Name	Type	Description
cluster	<a href="#">cluster</a>	
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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
volume	<a href="#">volume</a>	

#### svm

##### Origin volume SVM

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
flexcaches	array[ <a href="#">flexcache_relationship</a> ]	

Name	Type	Description
name	string	Origin volume name
svm	<a href="#">svm</a>	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[ <a href="#">error_arguments</a> ]	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}

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 use filter and retrieve only the required fields. See [DOC Requesting specific fields](#) to learn more.

- `flexcaches.ip_address` - IP address of FlexCache.
- `flexcaches.size` - Physical size 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
<a href="#">_links</a>	<a href="#">_links</a>	
flexcaches	array[ <a href="#">flexcache_relationship</a> ]	
name	string	Origin volume name
svm	<a href="#">svm</a>	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": "2018-06-04 19:00:00 UTC",
    "ip_address": "10.10.10.7",
    "size": 0,
    "state": "error",
    "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	<a href="#">href</a>	

cluster

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li></ul>

## flexcache\_relationship

Name	Type	Description
cluster	<a href="#">cluster</a>	
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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
volume	<a href="#">volume</a>	

## svm

## Origin volume SVM

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
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

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
serial_number	string	query	False	Filter by serial_number
description	string	query	False	Filter by description
error.corrective_action	string	query	False	Filter by error.corrective_action
error.message	string	query	False	Filter by error.message
mac_address	string	query	False	Filter by mac_address
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
state	string	query	False	Filter by state
name	string	query	False	Filter by name
speed	number	query	False	Filter by speed
part_number	string	query	False	Filter by part_number
wwn	string	query	False	Filter by wwn
cable.identifier	string	query	False	Filter by cable.identifier



Name	Type	In	Required	Description
cable.serial_number	string	query	False	Filter by cable.serial_number
cable.length	string	query	False	Filter by cable.length
cable.part_number	string	query	False	Filter by cable.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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
num_records	integer	Number of records
records	array[ <a href="#">storage_port</a> ]	

### Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "cable": {
      "identifier": "500a0980000b6c3f-50000d1703544b80",
      "length": "2m",
      "part_number": "112-00431+A0",
      "serial_number": 616930439
    },
    "description": "SAS Host Adapter 2a (PMC-Sierra PM8072 rev. C)",
    "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

href

Name	Type	Description
href	string	

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

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	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

storage\_port

Name	Type	Description
board_name	string	
cable	<a href="#">cable</a>	
description	string	
error	<a href="#">error</a>	
mac_address	string	
name	string	
node	<a href="#">node</a>	
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[ <a href="#">error_arguments</a> ]	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}

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	<a href="#">cable</a>	
description	string	
error	<a href="#">error</a>	
mac_address	string	
name	string	
node	<a href="#">node</a>	
part_number	string	
serial_number	string	
speed	number	Operational port speed in Gbps
state	string	
wwn	string	World Wide Name

## Example response

```
{
  "cable": {
    "identifier": "500a0980000b6c3f-50000d1703544b80",
    "length": "2m",
    "part_number": "112-00431+A0",
    "serial_number": 616930439
  },
  "description": "SAS Host Adapter 2a (PMC-Sierra PM8072 rev. C)",
  "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	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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 define 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 e.g. 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 would be limited to 1000 IOPS. If "capacity\_shared" is set to false then, each storage object will have it's SLOs enforced individually. For e.g. In the previous case if the same policy were 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 would provide 5000 IOPS per GB of allocated space for a storage object with a peak of 6000 IOPS. Minimum IOPS regardless of allocated space would be 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

Retrieves a collection of QoS policies.

Learn more

- [DOC /storage/qos/policies](#)

### Parameters

Name	Type	In	Required	Description
adaptive.expected_iops	integer	query	False	Filter by adaptive.expected_iops
adaptive.absolute_min_iops	integer	query	False	Filter by adaptive.absolute_min_iops
adaptive.peak_iops	integer	query	False	Filter by adaptive.peak_iops
uuid	string	query	False	Filter by uuid
name	string	query	False	Filter by name
object_count	integer	query	False	Filter by object_count
fixed.max_throughput_iops	integer	query	False	Filter by fixed.max_throughput_iops
fixed.capacity_shared	boolean	query	False	Filter by fixed.capacity_shared

Name	Type	In	Required	Description
fixed.min_throughput_iops	integer	query	False	Filter by fixed.min_throughput_iops
fixed.max_throughput_mbps	integer	query	False	Filter by fixed.max_throughput_mbps
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
error	<a href="#">error</a>	
num_records	integer	Number of records
records	array[ <a href="#">qos_policy</a> ]	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

\_links

Name	Type	Description
self	<a href="#">href</a>	

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.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

qos\_policy

Name	Type	Description
_links	<a href="#">_links</a>	
adaptive	<a href="#">adaptive</a>	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	<a href="#">fixed</a>	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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

## Create a QoS policy

POST /storage/qos/policies

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`

### Learn more

- [DOC /storage/qos/policies](#)

### Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	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.

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
adaptive	<a href="#">adaptive</a>	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	<a href="#">fixed</a>	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.

Name	Type	Description
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	

### Example request

```

{
  "_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"
}

```

### 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
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.
8454154	The name specified for creating conflicts with an existing QoS policy name.
8454147	The maximum limit for QoS policies has been reached.
8454273	Invalid value for an adaptive field. Value should be non-zero.
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.

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	<a href="#">href</a>	

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.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

qos\_policy

Name	Type	Description
_links	<a href="#">_links</a>	
adaptive	<a href="#">adaptive</a>	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
fixed	<a href="#">fixed</a>	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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

#### job\_link

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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/{policy.uuid}
```

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`

### Learn more

- [DOC /storage/qos/policies](#)

### Parameters

Name	Type	In	Required	Description
policy.uuid	string	path	True	
return_timeout	integer	query	False	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.

### 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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/{policy.uuid}

Retrieves a specific QoS policy.

### Related ONTAP commands

- `qos policy-group show`
- `qos adaptive-policy-group show`

### Learn more

- [DOC /storage/qos/policies](#)

### Parameters

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

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
adaptive	<a href="#">adaptive</a>	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	<a href="#">fixed</a>	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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.

Name	Type	Description
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	<a href="#">href</a>	

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.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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/{policy.uuid}

Update a specific QoS policy.

### Related ONTAP commands

- `qos policy-group modify`
- `qos adaptive-policy-group modify`

### Learn more

- [DOC /storage/qos/policies](#)

### Parameters

Name	Type	In	Required	Description
policy.uuid	string	path	True	

Name	Type	In	Required	Description
return_timeout	integer	query	False	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.

### Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
adaptive	<a href="#">adaptive</a>	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	<a href="#">fixed</a>	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.

Name	Type	Description
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	

### Example request

```

{
  "_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"
}

```

### 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
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.
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.
8454154	The name specified for creating conflicts with an existing QoS policy name.
8454147	The maximum limit for QoS policies has been reached.
8454273	Invalid value for an adaptive field. Value should be non-zero.
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.
8454286	Modifications on these cluster scoped preset policies is prohibited.

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	<a href="#">href</a>	

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.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

qos\_policy

Name	Type	Description
_links	<a href="#">_links</a>	
adaptive	<a href="#">adaptive</a>	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
fixed	<a href="#">fixed</a>	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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

#### job\_link

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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 or a FlexGroup volume.

#### Qtree APIs

The following APIs are used to create, retrieve, modify, and delete qtrees.

&ndash; POST /api/storage/qtrees

&ndash; GET /api/storage/qtrees

&ndash; GET /api/storage/qtrees/{volume-uuid}/{qtree-id}

&ndash; PATCH /api/storage/qtrees/{volume-uuid}/{qtree-id}

&ndash; 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 with a given security style, UNIX permissions, and an export 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"
}
}

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1"
      },
      "volume": {
        "name": "fv"
      },
      "name": "qt1",
      "security_style": "unix",
      "unix_permissions": 744,
      "export_policy": {
        "name": "default"
      },
      "_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=svml&volume.name=fv" -H 'accept:
application/hal+json'

# The response
{
  "records": [
    {
      "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"
          }
        }
      },
      "id": 0,
      "name": "",
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/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"
            }
        }
    },
    "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"
      }
    }
  },
  "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"
      }
    }
  },
  "_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": "expl"
  }
}
```

---

### 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

Retrieves qtrees configured for all FlexVols 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.

### Related ONTAP commands

- `qtree show`

### Learn more

- [DOC /storage/qtrees](#)

### Parameters

Name	Type	In	Required	Description
export_policy.name	string	query	False	Filter by export_policy.name
export_policy.id	integer	query	False	Filter by export_policy.id
name	string	query	False	Filter by name
security_style	string	query	False	Filter by security_style
volume.name	string	query	False	Filter by volume.name
volume.uuid	string	query	False	Filter by volume.uuid
svm.uuid	string	query	False	Filter by svm.uuid

Name	Type	In	Required	Description
svm.name	string	query	False	Filter by svm.name
path	string	query	False	Filter by path
unix_permissions	integer	query	False	Filter by unix_permissions
id	integer	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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
num_records	integer	Number of records
records	array[ <a href="#">qtree</a> ]	

## 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,
    "path": "/volume3/qtrees1",
    "security_style": "unix",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "unix_permissions": 493,
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

export\_policy

Export Policy

Name	Type	Description
_links	<a href="#">_links</a>	
id	integer	
name	string	

svm

Required in POST

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Required in POST

Name	Type	Description
_links	<a href="#">_links</a>	

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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

### 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	<a href="#">_links</a>	
export_policy	<a href="#">export_policy</a>	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.
security_style	string	Security style. Valid in POST or PATCH.
svm	<a href="#">svm</a>	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	<a href="#">volume</a>	Required in POST

### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

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`

## Learn more

- [DOC /storage/qtrees](#)

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
export_policy	<a href="#">export_policy</a>	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.
security_style	string	Security style. Valid in POST or PATCH.
svm	<a href="#">svm</a>	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	<a href="#">volume</a>	Required in POST

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "id": 1,
  "path": "/volume3/qtreen1",
  "security_style": "unix",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_permissions": 493,
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

## 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
2621462	The specified SVM does not exist.
917927	The specified volume was not found.
2621706	The specified <code>svm.uuid</code> and <code>svm.name</code> do not refer to the same SVM.
918236	The specified <code>volume.uuid</code> and <code>volume.name</code> refer to different volumes.
2621707	No SVM was specified. Either <code>svm.name</code> or <code>svm.uuid</code> must be supplied.
918232	Either <code>volume.name</code> or <code>volume.uuid</code> must be provided.
5242953	Qtree name must be provided.
5242952	Export Policy ID specified is invalid.
5242951	Export Policy supplied does not belong to the specified Export Policy ID.

Error Code	Description
5242886	Failed to create qtree.

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	<a href="#">href</a>	

export\_policy

Export Policy

Name	Type	Description
_links	<a href="#">_links</a>	
id	integer	
name	string	

svm

Required in POST

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Required in POST

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the volume.



Name	Type	Description
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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

### 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	<a href="#">_links</a>	
export_policy	<a href="#">export_policy</a>	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.
security_style	string	Security style. Valid in POST or PATCH.
svm	<a href="#">svm</a>	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	<a href="#">volume</a>	Required in POST

### job\_link

Name	Type	Description
_links	<a href="#">_links</a>	

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[ <a href="#">error_arguments</a> ]	Message arguments
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}

Deletes a qtree.

### Related ONTAP commands

- `qtree delete`

### Learn more

- [DOC /storage/qtrees](#)

### Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	
id	string	path	True	

## 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

Retrieves properties for a specific qtree identified by the `volume.uuid` and the `id` in the api path.

### Related ONTAP commands

- `qtree show`

### Learn more

- [DOC /storage/qtrees](#)

### 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>	<a href="#">_links</a>	
<code>export_policy</code>	<a href="#">export_policy</a>	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>security_style</code>	string	Security style. Valid in POST or PATCH.
<code>svm</code>	<a href="#">svm</a>	Required in POST

Name	Type	Description
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,
  "path": "/volume3/qtreen1",
  "security_style": "unix",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_permissions": 493,
  "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	<a href="#">href</a>	

export\_policy

Export Policy

Name	Type	Description
_links	<a href="#">_links</a>	
id	integer	
name	string	

svm

Required in POST

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Required in POST

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the volume.

Name	Type	Description
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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Update properties for a qtree

PATCH /storage/qtrees/{volume.uuid}/{id}

Updates properties for a specific qtree.

#### Related ONTAP commands

- `qtree modify`
- `qtree rename`

#### Learn more

- [DOC /storage/qtrees](#)

## Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
id	string	path	True	Qtree ID

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
export_policy	<a href="#">export_policy</a>	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.
security_style	string	Security style. Valid in POST or PATCH.
svm	<a href="#">svm</a>	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	<a href="#">volume</a>	Required in POST

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "id": 1,
  "path": "/volume3/qtreen1",
  "security_style": "unix",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_permissions": 493,
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

## 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.
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.

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	<a href="#">href</a>	

export\_policy

Export Policy

Name	Type	Description
_links	<a href="#">_links</a>	
id	integer	
name	string	

svm

Required in POST

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Required in POST

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the volume.



Name	Type	Description
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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

### 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	<a href="#">_links</a>	
export_policy	<a href="#">export_policy</a>	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.
security_style	string	Security style. Valid in POST or PATCH.
svm	<a href="#">svm</a>	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	<a href="#">volume</a>	Required in POST

### job\_link

Name	Type	Description
_links	<a href="#">_links</a>	

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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
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.

&ndash; POST /api/storage/quota/rules

&ndash; GET /api/storage/quota/rules

&ndash; GET /api/storage/quota/rules/{rule-uuid}

&ndash; PATCH /api/storage/quota/rules/{rule-uuid}

&ndash; 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.

&ndash; PATCH /api/storage/volumes/{volume-uuid} -d '{"quota.enabled":"true"}

&ndash; PATCH /api/storage/volumes/{volume-uuid} -d '{"quota.enabled":"false"}

&ndash; 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.

&ndash; GET /api/storage/quota/reports

&ndash; 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.

&ndash; PATCH /api/storage/quota/rules/{rule-uuid}<br> The changed limits in the filesystem can be confirmed using the quota report REST API.

&ndash; 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.

&ndash; GET /api/storage/quota/reports

&ndash; 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 FlexVols 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": "svml",
        "_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": "svml",
    "_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": "svml",
  "_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": "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": 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": "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": 1153202979583557632,
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/1153202979583557632"
      }
    }
  }
}

```

```

},
{
  "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": 2305843013508661248,
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/2305843013508661248"
    }
  }
},
{
  "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": 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": "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": 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 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 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": "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": 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 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"
  }
}
}
}

```

## Retrieve the quota report records for all FlexVol and FlexGroup volumes

GET /storage/quota/reports

Retrieves the quota report records for all FlexVols and FlexGroup volumes.

### Related ONTAP commands

- `quota report`

### Learn more

- [DOC /storage/quota/reports](#)

### Parameters

Name	Type	In	Required	Description
qtree.name	string	query	False	Filter by qtree.name
qtree.id	integer	query	False	Filter by qtree.id
index	integer	query	False	Filter by index

Name	Type	In	Required	Description
files.hard_limit	integer	query	False	Filter by files.hard_limit
files.soft_limit	integer	query	False	Filter by files.soft_limit
files.used.soft_limit_percent	string	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	string	query	False	Filter by files.used.hard_limit_percent
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
type	string	query	False	Filter by type
space.soft_limit	integer	query	False	Filter by space.soft_limit
space.hard_limit	integer	query	False	Filter by space.hard_limit
space.used.soft_limit_percent	string	query	False	Filter by space.used.soft_limit_percent
space.used.hard_limit_percent	string	query	False	Filter by space.used.hard_limit_percent
space.used.total	integer	query	False	Filter by space.used.total
volume.name	string	query	False	Filter by volume.name
volume.uuid	string	query	False	Filter by volume.uuid

Name	Type	In	Required	Description
users.id	string	query	False	Filter by users.id
users.name	string	query	False	Filter by users.name
group.name	string	query	False	Filter by group.name
group.id	string	query	False	Filter by group.id
specifier	string	query	False	Filter by specifier
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">quota_report</a> ]	

## 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": "string",
        "soft_limit_percent": "string",
        "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": "string",
        "soft_limit_percent": "string",

```



```

    "total": 0
  }
},
"specifier": "string",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"type": "tree",
"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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

used

Name	Type	Description
hard_limit_percent	string	Total files used as a percentage of file hard limit
soft_limit_percent	string	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	<a href="#">used</a>	

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
<a href="#">_links</a>	<a href="#">_links</a>	
id	integer	The unique identifier for a qtree.
name	string	The name of the qtree.

### used

Name	Type	Description
hard_limit_percent	string	Total space used as a percentage of space hard limit
soft_limit_percent	string	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	<a href="#">used</a>	

### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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
<a href="#">_links</a>	<a href="#">_links</a>	
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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

#### quota\_report

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
files	<a href="#">files</a>	
group	<a href="#">group</a>	<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	SVM, applies only to SVM-scoped objects.
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[ <a href="#">error_arguments</a> ]	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}

Retrieves a specific quota report record.

### Related ONTAP commands

- `quota report`

### Learn more

- [DOC /storage/quota/reports](#)

### 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	<a href="#">_links</a>	
files	<a href="#">files</a>	
group	<a href="#">group</a>	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	<a href="#">qtree</a>	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	<a href="#">space</a>	
specifier	string	Quota specifier
svm	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
type	string	Quota type associated with the quota record.



Name	Type	Description
users	array[users]	<p>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.</p>
volume	volume	

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "files": {
    "hard_limit": 0,
    "soft_limit": 0,
    "used": {
      "hard_limit_percent": "string",
      "soft_limit_percent": "string",
      "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": "string",
      "soft_limit_percent": "string",
      "total": 0
    }
  },
  "specifier": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  }
}
```

```

    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "tree",
  "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	<a href="#">href</a>	

used

Name	Type	Description
hard_limit_percent	string	Total files used as a percentage of file hard limit
soft_limit_percent	string	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	<a href="#">used</a>	

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	<a href="#">_links</a>	
id	integer	The unique identifier for a qtree.
name	string	The name of the qtree.

used

Name	Type	Description
hard_limit_percent	string	Total space used as a percentage of space hard limit
soft_limit_percent	string	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	<a href="#">used</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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
<a href="#">_links</a>	<a href="#">_links</a>	
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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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 FlexVols or FlexGroup volumes. Each quota rule has a type. The type can be "user", "group", or "tree".

&ndash; User rules must have the user property and qtree property.

&ndash; Group rules must have the group property and qtree property.

&ndash; 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.

&ndash; POST /api/storage/quota/rules

&ndash; GET /api/storage/quota/rules

&ndash; GET /api/storage/quota/rules/{rule-uuid}

&ndash; PATCH /api/storage/quota/rules/{rule-uuid}

&ndash; 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 FlexVols 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": "svml",
```



```

    "_links": {
      "self": {
        "href": "/api/svm/svms/038545f8-9ff8-11e8-bce6-005056a73bed"
      }
    },
    "volume": {
      "uuid": "ab3df793-0f02-43c6-9514-4f142fc8cc92",
      "name": "voll1",
      "_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": "voll1",
      "_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/{rule.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/c1b64eea-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/c1b64eea-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": "voll1",
    "_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": "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": "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": "voll1",
    "_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 quote policy rule.

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": "voll1"
},
"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": "svml1"
      },
      "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": "voll1"
  },
  "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": "voll1"
  },
  "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/{rule.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/{rule.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

Retrieves quota policy rules configured for all FlexVols and FlexGroup volumes.

### Related ONTAP commands

- `quota policy rule show`

### Learn more

- [DOC /storage/quota/rules](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	query	False	Filter by uuid
users.id	string	query	False	Filter by users.id
users.name	string	query	False	Filter by users.name
group.name	string	query	False	Filter by group.name
group.id	string	query	False	Filter by group.id
qtree.name	string	query	False	Filter by qtree.name
qtree.id	integer	query	False	Filter by qtree.id
user_mapping	boolean	query	False	Filter by user_mapping
files.soft_limit	integer	query	False	Filter by files.soft_limit
files.hard_limit	integer	query	False	Filter by files.hard_limit
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
type	string	query	False	Filter by type
space.soft_limit	integer	query	False	Filter by space.soft_limit



Name	Type	In	Required	Description
space.hard_limit	integer	query	False	Filter by space.hard_limit
volume.name	string	query	False	Filter by volume.name
volume.uuid	string	query	False	Filter by volume.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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	

Name	Type	Description
num_records	integer	Number of records
records	array[ <a href="#">quota_rule</a> ]	

## 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": "tree",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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	<a href="#">_links</a>	
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.
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.

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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
<a href="#">_links</a>	<a href="#">_links</a>	
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"><li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li></ul>

## quota\_rule

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
files	<a href="#">files</a>	
group	<a href="#">group</a>	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	SVM, applies only to SVM-scoped objects.
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	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

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.soft_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`

## Learn more

- [DOC /storage/quota/rules](#)

## Request Body

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>files</code>	<code>files</code>	
<code>group</code>	<code>group</code>	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.
<code>qtree</code>	<code>qtree</code>	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
space	space	
svm	svm	SVM, applies only to SVM-scoped objects.
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**

A large, empty rectangular box with a thin, dashed border, occupying most of the page. It is intended for an example request.

```

{
  "_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": "tree",
  "users": {
    "id": "string",
    "name": "string"
  },
  "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volumel",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}

```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

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

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
2621462	The specified SVM does not exist.
917927	The specified volume was not found.
2621706	The specified <code>svm.uuid</code> and <code>svm.name</code> do not refer to the same SVM.
918236	The specified <code>volume.uuid</code> and <code>volume.name</code> refer to different volumes.
2621707	No SVM was specified. Either <code>svm.name</code> or <code>svm.uuid</code> must be supplied.
918232	Either <code>volume.name</code> or <code>volume.uuid</code> must be provided.
5308549	Volume does not belong to SVM.
5308552	Failed to get default quota policy name for SVM.



Error Code	Description
5308562	users is a required input for creating a user rule and group is not allowed.
5308563	group is a required input for creating a group rule and users is not allowed.
5308564	qtree.name is a required input for creating a tree rule and users and group are not allowed.
5308566	Only one of name or id is allowed for group.
5308565	Only one of name or id is allowed for each entry in the users array.
5308561	Failed to obtain volume quota state or invalid quota state obtained for volume.
5308559	Cannot create the quota rule for volume because it is a default quota rule and quotas are enabled on the volume..
5308560	Cannot create the explicit quota rule for volume because it does not have a default quota rule and quotas are enabled on the volume.
5308568	Quota policy rule create operation succeeded, but quota resize failed due to internal error.
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.

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	<a href="#">href</a>	

### 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
<a href="#">_links</a>	<a href="#">_links</a>	
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.
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.

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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
<a href="#">_links</a>	<a href="#">_links</a>	

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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

quota\_rule

Name	Type	Description
_links	<a href="#">_links</a>	
files	<a href="#">files</a>	
group	<a href="#">group</a>	<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	SVM, applies only to SVM-scoped objects.
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
_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[ <a href="#">error_arguments</a> ]	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/{rule.uuid}
```

Deletes a quota policy rule.

### Related ONTAP commands

- `quota policy rule delete`

### Learn more

- [DOC /storage/quota/rules](#)

### Parameters

Name	Type	In	Required	Description
rule.uuid	string	path	True	

### Response

```
Status: 202, Accepted
```

Name	Type	Description
job	<a href="#">job_link</a>	



## 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.
5308547	Failed to delete the quota policy rule.
5308567	Quota policy rule delete operation succeeded, but quota resize failed due to internal 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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/{rule.uuid}

Retrieves properties for a specific quota policy rule.

#### Related ONTAP commands

- `quota policy rule show`

#### Learn more

- [DOC /storage/quota/rules](#)

#### Parameters

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

#### Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
<a href="#">files</a>	<a href="#">files</a>	
<a href="#">group</a>	<a href="#">group</a>	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	SVM, applies only to SVM-scoped objects.
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]	<p>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.</p>
uuid	string	<p>Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.</p>
volume	volume	

**Example response**

A large, empty rectangular box with a thin, dashed border, occupying most of the page. It is intended for an example response.

```

{
  "_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": "tree",
  "users": {
    "id": "string",
    "name": "string"
  },
  "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volumel",
    "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.

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	<a href="#">href</a>	

### 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
<a href="#">_links</a>	<a href="#">_links</a>	
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.
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.

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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
<a href="#">_links</a>	<a href="#">_links</a>	

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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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/{rule.uuid}

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`

## Learn more

- [DOC /storage/quota/rules](#)

## Parameters

Name	Type	In	Required	Description
rule.uuid	string	path	True	Rule UUID

## Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>files</code>	<a href="#">files</a>	
<code>group</code>	<a href="#">group</a>	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	SVM, applies only to SVM-scoped objects.
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]	<p>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.</p>
uuid	string	<p>Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.</p>
volume	volume	

**Example request**





```

{
  "_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": "tree",
  "users": {
    "id": "string",
    "name": "string"
  },
  "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volumel",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}

```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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.
5308546	Failed to modify the quota policy rule.
5308567	Quota policy rule modify operation succeeded, but quota resize failed due to internal error.
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.

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

### 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	<a href="#">_links</a>	
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.
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.

#### svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	

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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

quota\_rule

Name	Type	Description
_links	<a href="#">_links</a>	
files	<a href="#">files</a>	
group	<a href="#">group</a>	<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	SVM, applies only to SVM-scoped objects.
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
_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[ <a href="#">error_arguments</a> ]	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,  
  "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,
    "wnn": "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,
  "wwn": "500A098000C9EDBF",
  "cable": {
    "identifier": "500a098000d5c4bf-500a098000c9edbf",
    "part_number": "112-00176+A0",
    "length": "0.5-1.0m",
    "serial_number": "APF133917610YT"
  },
  "remote": {
    "wwn": "500A098000D5C4BF",
    "phy": "00"
  }
},
{
  "id": 2,
  "module_id": "b",
  "designator": "square",
  "state": "connected",
  "internal": false,
  "wwn": "500A098004F208BF",
  "cable": {
    "identifier": "5001086000702648-500a098004f208bf",
    "part_number": "112-00430+A0",
    "length": "2m",
    "serial_number": "APF16510229540"
  },
  "remote": {
    "wwn": "5001086000702648",
    "phy": "08"
  }
},

```

```

{
  "id": 3,
  "module_id": "b",
  "designator": "circle",
  "state": "connected",
  "internal": false,
  "wwn": "500A098004F208BF",
  "cable": {
    "identifier": "500a0980062ba33f-500a098004f208bf",
    "part_number": "112-00176+20",
    "length": "0.5-1.0m",
    "serial_number": "832210017"
  },
  "remote": {
    "wwn": "500A0980062BA33F",
    "phy": "00"
  }
},
]
"_links": {
  "self": {
    "href": "/api/storage/shelves/3109174803597886800"
  }
}
}

```

## Retrieve shelves

GET /storage/shelves

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
paths.name	string	query	False	Filter by paths.name
id	string	query	False	Filter by id
serial_number	string	query	False	Filter by serial_number
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
connection_type	string	query	False	Filter by connection_type
module_type	string	query	False	Filter by module_type
drawers.disk_count	integer	query	False	Filter by drawers.disk_count
drawers.part_number	string	query	False	Filter by drawers.part_number
drawers.error	string	query	False	Filter by drawers.error

Name	Type	In	Required	Description
drawers.id	integer	query	False	Filter by drawers.id
drawers.closed	boolean	query	False	Filter by drawers.closed
drawers.state	string	query	False	Filter by drawers.state
drawers.serial_number	string	query	False	Filter by drawers.serial_number
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
disk_count	integer	query	False	Filter by disk_count
model	string	query	False	Filter by model
name	string	query	False	Filter by name
state	string	query	False	Filter by state
ports.state	string	query	False	Filter by ports.state
ports.module_id	string	query	False	Filter by ports.module_id
ports.cable.serial_number	string	query	False	Filter by ports.cable.serial_number
ports.cable.identifier	string	query	False	Filter by ports.cable.identifier
ports.cable.part_number	string	query	False	Filter by ports.cable.part_number

Name	Type	In	Required	Description
ports.cable.length	string	query	False	Filter by ports.cable.length
ports.wwn	string	query	False	Filter by ports.wwn
ports.internal	boolean	query	False	Filter by ports.internal
ports.designator	string	query	False	Filter by ports.designator
ports.id	integer	query	False	Filter by ports.id
ports.remote.phy	string	query	False	Filter by ports.remote.phy
ports.remote.port	string	query	False	Filter by ports.remote.port
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.wwn	string	query	False	Filter by ports.remote.wwn
ports.mac_address	string	query	False	Filter by ports.mac_address
internal	boolean	query	False	Filter by internal
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.

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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">shelf</a> ]	

## 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,
      "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,
  "module_id": "a",
  "remote": {
    "phy": 12,
    "wwn": "50000D1703544B80"
  },
  "state": "connected",
  "wwn": "500A0980000B6C3F"
},
"serial_number": "SHFMS1514000895",
"state": "ok",
"uid": 7777841915827391056
}
}

```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

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	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

paths

Storage port

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

remote

Name	Type	Description
chassis	string	
mac_address	string	
phy	string	
port	string	
wwn	string	

ports

Name	Type	Description
cable	<a href="#">cable</a>	
designator	string	
id	integer	
internal	boolean	
mac_address	string	
module_id	string	
remote	<a href="#">remote</a>	
state	string	
wwn	string	

#### shelf

Name	Type	Description
bays	array[ <a href="#">bays</a> ]	
connection_type	string	
disk_count	integer	
drawers	array[ <a href="#">drawers</a> ]	
frus	array[ <a href="#">frus</a> ]	
id	string	
internal	boolean	
model	string	
module_type	string	
name	string	
paths	array[ <a href="#">paths</a> ]	
ports	array[ <a href="#">ports</a> ]	
serial_number	string	
state	string	
uid	string	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}

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[ <a href="#">bays</a> ]	
connection_type	string	

Name	Type	Description
disk_count	integer	
drawers	array[drawers]	
frus	array[frus]	
id	string	
internal	boolean	
model	string	
module_type	string	
name	string	
paths	array[paths]	
ports	array[ports]	
serial_number	string	
state	string	
uid	string	

## Example response

```
{
  "bays": {
    "id": 0,
    "state": "ok",
    "type": "single_disk"
  },
  "connection_type": "sas",
  "disk_count": 12,
  "drawers": {
    "disk_count": 12,
    "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,
  "module_id": "a",
  "remote": {
    "phy": 12,
    "wwn": "50000D1703544B80"
  },
  "state": "connected",
  "wwn": "500A0980000B6C3F"
},
"serial_number": "SHFMS1514000895",
"state": "ok",
"uid": 7777841915827391056
}

```

#### 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	<a href="#">href</a>	

## node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

## paths

### Storage port

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
node	<a href="#">node</a>	

## cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

## remote

Name	Type	Description
chassis	string	
mac_address	string	
phy	string	
port	string	
wwn	string	

## ports

Name	Type	Description
cable	<a href="#">cable</a>	
designator	string	
id	integer	
internal	boolean	
mac_address	string	

Name	Type	Description
module_id	string	
remote	<a href="#">remote</a>	
state	string	
wwn	string	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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, which are "default", "default-1weekly" and "none" policies.

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. So once the Snapshot copy count reaches the retention count, on the next create the oldest Snapshot copy is deleted.

## Snapshot copy policy APIs

The following APIs are used to perform operations related to Snapshot copy policy information:

&ndash; POST /api/storage/snapshot\_policies

&ndash; GET /api/storage/snapshot\_policies

&ndash; GET /api/storage/snapshot\_policies/{uuid}

&ndash; PATCH /api/storage/snapshot\_policies/{uuid}

&ndash; 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

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
uuid	string	query	False	Filter by uuid
comment	string	query	False	Filter by comment
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
enabled	boolean	query	False	Filter by enabled
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
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.

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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">snapshot_policy</a> ]	

## 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"
      }
    },
    "enabled": 1,
    "name": "default",
    "scope": "svm",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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	<a href="#">schedule</a>	
snapmirror_label	string	Label for SnapMirror operations

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[ <a href="#">copies</a> ]	
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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error



Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}

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	

#### 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.

Error Code	Description
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[ <a href="#">error_arguments</a> ]	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}

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 ID
fields	array[string]	query	False	Specify the fields to return.

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[ <a href="#">copies</a> ]	
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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "copies": {
    "prefix": "string",
    "schedule": {
      "name": "hourly"
    }
  },
  "enabled": 1,
  "name": "default",
  "scope": "svm",
  "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	<a href="#">href</a>	

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	<a href="#">schedule</a>	
snapmirror_label	string	Label for SnapMirror operations

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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 policy UUID



## Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[ <a href="#">copies</a> ]	
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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "copies": {
    "prefix": "string",
    "schedule": {
      "name": "hourly"
    }
  },
  "enabled": 1,
  "name": "default",
  "scope": "svm",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "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.

Name	Type	Description
error	<a href="#">error</a>	

### 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	<a href="#">href</a>	

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	<a href="#">schedule</a>	
snamirror_label	string	Label for SnapMirror operations

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.

Name	Type	Description
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	<a href="#">_links</a>	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[ <a href="#">copies</a> ]	
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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

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](#)

### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>comment</code>	string	A comment associated with the Snapshot copy policy.

Name	Type	Description
copies	array[ <a href="#">copies</a> ]	
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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "copies": {
    "prefix": "string",
    "schedule": {
      "name": "hourly"
    }
  },
  "enabled": 1,
  "name": "default",
  "scope": "svm",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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.
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.

Name	Type	Description
error	<a href="#">error</a>	

### 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	<a href="#">href</a>	

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	<a href="#">schedule</a>	
snapmirror_label	string	Label for SnapMirror operations

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.

Name	Type	Description
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	<a href="#">_links</a>	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[ <a href="#">copies</a> ]	
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	<a href="#">svm</a>	SVM, applies only to SVM-scoped objects.
uuid	string	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code

Name	Type	Description
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.

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 will be enabled by default on volume creation. Options such as "background/inline/both" will be treated as both, which means both background and inline will be enabled for any efficiency option. The option "none" will disable 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:

&ndash; PATCH /api/storage/volumes/{volume-uuid} -d '{"quota.enabled":"true"}

&ndash; PATCH /api/storage/volumes/{volume-uuid} -d '{"quota.enabled":"false"}

&ndash; GET /api/storage/volumes/{volume-uuid}/?fields=quota.state

## 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 an a node which 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 FlexVols and FlexGroup volumes:

&ndash; POST `/api/storage/volumes`

&ndash; GET `/api/storage/volumes`

&ndash; GET `/api/storage/volumes/{uuid}`

&ndash; PATCH `/api/storage/volumes/{uuid}`

&ndash; 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 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": "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"
      }
    }
  }
}
```

```
# 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": "voll1"}}, {"svm":{"name": "vs0"}, {"clone":
"is_flexclone":"true"}}'

# 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:

&ndash; RW, DP and LS volume

&ndash; FlexGroup volume

&ndash; FlexCache volume

&ndash; FlexClone volume

#### The following types of volumes are not reported:

&ndash; DEL volume

&ndash; TEMP volume

&ndash; Node Root volume



&ndash; System Vserver volume

&ndash; FlexGroup constituent

&ndash; FlexCache constituent

## Examples

### Retrieving the attributes of a volume

```
# 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 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",
```

```
"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
  }
},
"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": "2018-08-20T13:17:43+05:30",
  "expiry_time": "2018-10-30T05:30:00+05:30",
  "is_audit_log": false,
  "litigation_count": 0,
  "privileged_delete": "permanently_disabled",
  "type": "compliance",
  "retention": {
    "default": "30 years",
    "minimum": "0 years",
    "maximum": "30 years"
  }
},
"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"
      }
    }
  }
}
```

## 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/{volume.id} " -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

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 [DOC Requesting specific fields](#) to learn more.

- `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.footprint
- space.over\_provisioned
- space.metadata
- 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

#### Learn more

- [DOC /storage/volumes](#)

#### Parameters

Name	Type	In	Required	Description
nas.unix_permissions	integer	query	False	Filter by nas.unix_permissions
nas.export_policy.name	string	query	False	Filter by nas.export_policy.name
nas.export_policy.id	integer	query	False	Filter by nas.export_policy.id
nas.security_style	string	query	False	Filter by nas.security_style
nas.gid	integer	query	False	Filter by nas.gid
nas.path	string	query	False	Filter by nas.path
nas.uid	integer	query	False	Filter by nas.uid
error_state.has_bad_blocks	boolean	query	False	Filter by error_state.has_bad_blocks
error_state.is_inconsistent	boolean	query	False	Filter by error_state.is_inconsistent

Name	Type	In	Required	Description
guarantee.honored	boolean	query	False	Filter by guarantee.honored
guarantee.type	string	query	False	Filter by guarantee.type
flexcache_endpoint_type	string	query	False	Filter by flexcache_endpoint_type
snaplock.retention.default	string	query	False	Filter by snaplock.retention.default
snaplock.retention.maximum	string	query	False	Filter by snaplock.retention.maximum
snaplock.retention.minimum	string	query	False	Filter by snaplock.retention.minimum
snaplock.append_mode_enabled	boolean	query	False	Filter by snaplock.append_mode_enabled
snaplock.autocommit_period	string	query	False	Filter by snaplock.autocommit_period
snaplock.type	string	query	False	Filter by snaplock.type
snaplock.litigation_count	integer	query	False	Filter by snaplock.litigation_count
snaplock.privileged_delete	string	query	False	Filter by snaplock.privileged_delete
snaplock.compliance_clock_time	string	query	False	Filter by snaplock.compliance_clock_time

Name	Type	In	Required	Description
snaplock.is_audit_log	boolean	query	False	Filter by snaplock.is_audit_log
snaplock.expiry_time	string	query	False	Filter by snaplock.expiry_time
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.iops.total	integer	query	False	Filter by metric.iops.total
metric.iops.other	integer	query	False	Filter by metric.iops.other
metric.iops.read	integer	query	False	Filter by metric.iops.read
metric.iops.write	integer	query	False	Filter by metric.iops.write
metric.latency.total	integer	query	False	Filter by metric.latency.total
metric.latency.other	integer	query	False	Filter by metric.latency.other
metric.latency.read	integer	query	False	Filter by metric.latency.read
metric.latency.write	integer	query	False	Filter by metric.latency.write
metric.throughput.total	integer	query	False	Filter by metric.throughput.total

Name	Type	In	Required	Description
metric.throughput.ot her	integer	query	False	Filter by metric.throughput.ot her
metric.throughput.re ad	integer	query	False	Filter by metric.throughput.re ad
metric.throughput.wri te	integer	query	False	Filter by metric.throughput.wr ite
language	string	query	False	Filter by language
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
movement.percent_c omplete	string	query	False	Filter by movement.percent_ complete
movement.destinatio n_aggregate.uuid	string	query	False	Filter by movement.destinatio n_aggregate.uuid
movement.destinatio n_aggregate.name	string	query	False	Filter by movement.destinatio n_aggregate.name
movement.cutover_ window	integer	query	False	Filter by movement.cutover_ window
movement.state	string	query	False	Filter by movement.state
efficiency.dedupe	string	query	False	Filter by efficiency.dedupe
efficiency.cross_volu me_dedupe	string	query	False	Filter by efficiency.cross_volu me_dedupe



<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
efficiency.compaction	string	query	False	Filter by efficiency.compaction
efficiency.compression	string	query	False	Filter by efficiency.compression
space.footprint	integer	query	False	Filter by space.footprint
space.capacity_tier_footprint	integer	query	False	Filter by space.capacity_tier_footprint
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.metadata	integer	query	False	Filter by space.metadata
space.available	integer	query	False	Filter by space.available
space.block_storage_inactive_user_data	integer	query	False	Filter by space.block_storage_inactive_user_data
space.over_provisioned	integer	query	False	Filter by space.over_provisioned
space.size	integer	query	False	Filter by space.size
space.logical_space.enforcement	boolean	query	False	Filter by space.logical_space.enforcement
space.logical_space.available	integer	query	False	Filter by space.logical_space.available

Name	Type	In	Required	Description
space.logical_space.reporting	boolean	query	False	Filter by space.logical_space.reporting
space.logical_space.used_by_afs	integer	query	False	Filter by space.logical_space.used_by_afs
space.used	integer	query	False	Filter by space.used
state	string	query	False	Filter by state
uuid	string	query	False	Filter by uuid
application.name	string	query	False	Filter by application.name
application.uuid	string	query	False	Filter by application.uuid
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
encryption.state	string	query	False	Filter by encryption.state
encryption.type	string	query	False	Filter by encryption.type
encryption.key_id	string	query	False	Filter by encryption.key_id
encryption.enabled	boolean	query	False	Filter by encryption.enabled
statistics.iops_raw.total	integer	query	False	Filter by statistics.iops_raw.total

Name	Type	In	Required	Description
statistics.iops_raw.other	integer	query	False	Filter by statistics.iops_raw.other
statistics.iops_raw.read	integer	query	False	Filter by statistics.iops_raw.read
statistics.iops_raw.write	integer	query	False	Filter by statistics.iops_raw.write
statistics.throughput_raw.total	integer	query	False	Filter by statistics.throughput_raw.total
statistics.throughput_raw.other	integer	query	False	Filter by statistics.throughput_raw.other
statistics.throughput_raw.read	integer	query	False	Filter by statistics.throughput_raw.read
statistics.throughput_raw.write	integer	query	False	Filter by statistics.throughput_raw.write
statistics.timestamp	string	query	False	Filter by statistics.timestamp
statistics.status	string	query	False	Filter by statistics.status
statistics.latency_raw.total	integer	query	False	Filter by statistics.latency_raw.total
statistics.latency_raw.other	integer	query	False	Filter by statistics.latency_raw.other
statistics.latency_raw.read	integer	query	False	Filter by statistics.latency_raw.read

<b>Name</b>	<b>Type</b>	<b>In</b>	<b>Required</b>	<b>Description</b>
statistics.latency_raw.write	integer	query	False	Filter by statistics.latency_raw.write
size	integer	query	False	Filter by size
tiering.policy	string	query	False	Filter by tiering.policy
files.used	integer	query	False	Filter by files.used
files.maximum	integer	query	False	Filter by files.maximum
aggregates.uuid	string	query	False	Filter by aggregates.uuid
aggregates.name	string	query	False	Filter by aggregates.name
style	string	query	False	Filter by style
create_time	string	query	False	Filter by create_time
quota.state	string	query	False	Filter by quota.state
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.split_estimate	integer	query	False	Filter by clone.split_estimate
clone.parent_snapshot.name	string	query	False	Filter by clone.parent_snapshot.name
clone.parent_snapshot.uuid	string	query	False	Filter by clone.parent_snapshot.uuid

Name	Type	In	Required	Description
clone.is_flexclone	boolean	query	False	Filter by clone.is_flexclone
clone.parent_volume.name	string	query	False	Filter by clone.parent_volume.name
clone.parent_volume.uuid	string	query	False	Filter by clone.parent_volume.uuid
clone.split_initiated	boolean	query	False	Filter by clone.split_initiated
clone.split_complete_percent	integer	query	False	Filter by clone.split_complete_percent
qos.policy.uuid	string	query	False	Filter by qos.policy.uuid
qos.policy.min_throughput_iops	integer	query	False	Filter by qos.policy.min_throughput_iops
qos.policy.max_throughput_mbps	integer	query	False	Filter by qos.policy.max_throughput_mbps
qos.policy.name	string	query	False	Filter by qos.policy.name
qos.policy.max_throughput_iops	integer	query	False	Filter by qos.policy.max_throughput_iops
snapshot_policy.uuid	string	query	False	Filter by snapshot_policy.uuid
snapshot_policy.name	string	query	False	Filter by snapshot_policy.name
name	string	query	False	Filter by name

Name	Type	In	Required	Description
type	string	query	False	Filter by type
comment	string	query	False	Filter by comment
autosize.minimum	integer	query	False	Filter by autosize.minimum
autosize.mode	string	query	False	Filter by autosize.mode
autosize.grow_threshold	integer	query	False	Filter by autosize.grow_threshold
autosize.maximum	integer	query	False	Filter by autosize.maximum
autosize.shrink_threshold	integer	query	False	Filter by autosize.shrink_threshold
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.
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.

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	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">volume</a> ]	

## 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"
    },
    "application": {
      "name": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
    },
    "autosize": {
      "mode": "grow"
    },
    "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": "svml",
    "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
},
"comment": "string",
"create_time": "2018-06-04 19:00:00 UTC",
"efficiency": {
    "compaction": "inline",
    "compression": "inline",
    "cross_volume_dedupe": "inline",
    "dedupe": "inline"
},
"encryption": {
    "key_id": "string",
    "state": "encrypted",
    "status": {
        "code": "string",
        "message": "string"
    },
    },
    "type": "none"
},
"files": {
    "used": 0
},
"flexcache_endpoint_type": "none",
"guarantee": {
    "type": "volume"
},
"language": "ar",
"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": "2017-01-25 11:20:13 UTC"
},
"movement": {
  "cutover_window": 30,
  "destination_aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "percent_complete": "string",
  "state": "replicating",
  "tiering_policy": "all"
},
"name": "vol_cs_dept",
"nas": {
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,

```

```

    "name": "default"
  },
  "path": "/user/my_volume",
  "security_style": "mixed",
  "unix_permissions": 493
},
"qos": {
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput_iops": 2000,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"quota": {
  "state": "corrupt"
},
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04 19:00:00 UTC",
  "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"
},
"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,
  "logical_space": {
    "available": 0,
    "used_by_afs": 0
  },
  "metadata": 0,
  "over_provisioned": 0,
  "snapshot": {
    "used": 0
  },
  "used": 0
},
"state": "error",
"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": "2017-01-25 11:20:13 UTC"
},
"style": "flexvol",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}

```

```
    },
    "tiering": {
      "policy": "all"
    },
    "type": "rw",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

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 &dash; Volume automatically grows when the amount of used space is above the 'grow_threshold' value.  grow_shrink &dash; Volume grows or shrinks in response to the amount of space used.  off &dash; 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	<a href="#">_links</a>	
name	string	
uuid	string	

#### parent\_svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### parent\_volume

Name	Type	Description
_links	<a href="#">_links</a>	
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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

#### 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	<a href="#">snapshot_reference</a>	
parent_svm	<a href="#">parent_svm</a>	SVM, applies only to SVM-scoped objects.
parent_volume	<a href="#">parent_volume</a>	
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.

#### efficiency

Name	Type	Description
compaction	string	<p>The system can be enabled/disabled compaction.&lt;br&gt; inline &amp;dash;  Data will be compacted first and written to the volume.&lt;br&gt; none &amp;dash;  None&lt;br&gt; mixed &amp;dash;  Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.</p>
compression	string	<p>The system can be enabled/disabled compression.&lt;br&gt; inline &amp;dash;  Data will be compressed first and written to the volume.&lt;br&gt; background &amp;dash;  Data will be written to the volume and compressed later.&lt;br&gt; both &amp;dash;  Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run.&lt;br&gt; none &amp;dash;  None&lt;br&gt; mixed &amp;dash;  Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.</p>

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 &dash; Data will be cross volume deduped first and written to the volume.  background &dash; Data will be written to the volume and cross volume deduped later.  both &dash; 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 &dash; None  mixed &dash; 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 &dash; Data will be deduped first and written to the volume.  background &dash; Data will be written to the volume and deduped later.  both &dash; Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run.  none &dash; None  mixed &dash; Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.

status

Name	Type	Description
code	string	Encryption progress message code.
message	string	Encryption progress message.

encryption

Name	Type	Description
enabled	boolean	Encrypts an unencrypted volume. When set to 'true', a new key is generated and used to encrypt the given volume. The underlying SVM must be configured with the key manager.
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 &dash; The volume is completely encrypted.  encrypting &dash; Encryption operation is in progress.  partial &dash; 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 &dash; The volume is a plain-text one.
status	<a href="#">status</a>	
type	string	Volume encryption type.  none &dash; The volume is a plain-text one.  volume &dash; The volume is encrypted with volume key (NVE volume).  aggregate &dash; The volume is encrypted with aggregate key (NAE volume).

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 &dash; File system is inconsistent.  false &dash; File system in 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?

Name	Type	Description
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.
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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">iops</a>	The rate of I/O operations observed at the storage object.
latency	<a href="#">latency</a>	The round trip latency in microseconds observed at the storage object.

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 does not have the latest data.
throughput	<a href="#">throughput</a>	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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">destination_aggregate</a>	Aggregate
percent_complete	string	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. 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	<a href="#">_links</a>	
id	integer	
name	string	

nas

Name	Type	Description
export_policy	<a href="#">export_policy</a>	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. 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 &dash; Mixed-style security ntfs &dash; NTFS/Windows-style security unified &dash; Unified-style security, unified UNIX, NFS and CIFS permissions unix &dash; 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", "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
<a href="#">_links</a>	<a href="#">_links</a>	
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. Default is 15000 on AFF platforms and 10000 on all other platforms.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name, UUID and "min_throughput_iops" 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, UUID and "max_throughput_mbps" 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	<a href="#">policy</a>	When "min_throughput_iops", "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	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.

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&lt;num&gt;Y", "P&lt;num&gt;M", "P&lt;num&gt;D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT&lt;num&gt;H" and "PT&lt;num&gt;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	<p>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&lt;num&gt;Y", "P&lt;num&gt;M", "P&lt;num&gt;D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT&lt;num&gt;H" and "PT&lt;num&gt;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>

snaplock

Name	Type	Description
append_mode_enabled	boolean	<p>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.</p>

Name	Type	Description
autocommit_period	string	<p>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&lt;num&gt;Y", "P&lt;num&gt;M", "P&lt;num&gt;D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT&lt;num&gt;H" and "PT&lt;num&gt;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".</p>
compliance_clock_time	string	<p>This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.</p>
expiry_time	string	<p>Expiry time of the volume.</p>
is_audit_log	boolean	<p>Indicates if this volume has been configured as SnapLock audit log volume for the SVM .</p>
litigation_count	integer	<p>Litigation count indicates the number of active legal-holds on the volume.</p>



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	<a href="#">retention</a>	
type	string	The SnapLock type of the volume.   compliance &dash; 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 &dash; An administrator can delete a SnapLock Enterprise(SLE) volume.   non_snaplock &dash; Indicates the volume is non-snaplock.

#### snapshot\_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	<a href="#">_links</a>	
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.

Name	Type	Description
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.

Name	Type	Description
capacity_tier_footprint	integer	The space used by capacity tier for this volume in the aggregate, in bytes.
footprint	integer	Data and metadata used for this volume in the aggregate, in bytes.
logical_space	<a href="#">logical_space</a>	
metadata	integer	The space used by the total metadata in the volume, in bytes.
over_provisioned	integer	The amount of space not available for this volume in the aggregate, in bytes.
size	integer	Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes.
snapshot	<a href="#">snapshot</a>	
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 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.
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 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 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	<p>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".</p> <p>"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.</p> <p>"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.</p>
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

SVM containing the volume. Required on POST.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

tiering

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 &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.  auto &dash; This policy allows tiering of both snapshot and active file system user data to the cloud store  none &dash; Volume blocks will not be tiered to the cloud store.  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.

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
_links	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	Aggregate hosting the volume. Required on POST.
application	<a href="#">application</a>	
autosize	<a href="#">autosize</a>	
clone	<a href="#">clone</a>	
comment	string	A comment for the volume. Valid in POST or PATCH.

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	<a href="#">efficiency</a>	
encryption	<a href="#">encryption</a>	
error_state	<a href="#">error_state</a>	
files	<a href="#">files</a>	
flexcache_endpoint_type	string	FlexCache endpoint type. none &dash; The volume is neither a FlexCache nor origin of any FlexCache. cache &dash; The volume is a FlexCache volume. origin &dash; The volume is origin of a FlexCache volume.
guarantee	<a href="#">guarantee</a>	
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	<a href="#">metric</a>	Performance numbers, such as IOPS latency and throughput.



Name	Type	Description
movement	<a href="#">movement</a>	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	<a href="#">nas</a>	
qos	<a href="#">qos</a>	QoS information
quota	<a href="#">quota</a>	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. 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	<a href="#">snaplock</a>	
snapshot_policy	<a href="#">snapshot_policy</a>	This is a reference to the Snapshot copy policy.
space	<a href="#">space</a>	

Name	Type	Description
state	string	Volume state. A volume can only be brought online if it is offline. 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	<a href="#">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.
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. If "style" is specified, a volume of that type is created. That is, if style is "flexvol", a single aggregate must be specified. If style is "flexgroup", the system either uses the specified aggregates, or automatically provisions if no aggregates are specified. flexvol &dash; flexible volumes and FlexClone volumes flexgroup &dash; FlexGroups.
svm	<a href="#">svm</a>	SVM containing the volume. Required on POST.
tiering	<a href="#">tiering</a>	
type	string	Type of the volume. rw &dash; read-write volume. dp &dash; data-protection volume. ls &dash; 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	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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> <li>• readOnly: 1</li> </ul>

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

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

## Create a volume on an SVM and storage aggregates

POST /storage/volumes

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`

### Learn more

- [DOC /storage/volumes](#)

### Request Body

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>aggregates</code>	<code>array[aggregates]</code>	Aggregate hosting the volume. Required on POST.
<code>application</code>	<code>application</code>	

Name	Type	Description
autosize	<a href="#">autosize</a>	
clone	<a href="#">clone</a>	
comment	string	A comment for the volume. Valid in POST or PATCH.
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	<a href="#">efficiency</a>	
encryption	<a href="#">encryption</a>	
error_state	<a href="#">error_state</a>	
files	<a href="#">files</a>	
flexcache_endpoint_type	string	FlexCache endpoint type.  none &dash; The volume is neither a FlexCache nor origin of any FlexCache.  cache &dash; The volume is a FlexCache volume.  origin &dash; The volume is origin of a FlexCache volume.
guarantee	<a href="#">guarantee</a>	
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	<a href="#">metric</a>	Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
movement	<a href="#">movement</a>	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	<a href="#">nas</a>	
qos	<a href="#">qos</a>	QoS information
quota	<a href="#">quota</a>	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. 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	<a href="#">snaplock</a>	
snapshot_policy	<a href="#">snapshot_policy</a>	This is a reference to the Snapshot copy policy.
space	<a href="#">space</a>	

Name	Type	Description
state	string	Volume state. A volume can only be brought online if it is offline. 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	<a href="#">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.
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. If "style" is specified, a volume of that type is created. That is, if style is "flexvol", a single aggregate must be specified. If style is "flexgroup", the system either uses the specified aggregates, or automatically provisions if no aggregates are specified. flexvol &dash; flexible volumes and FlexClone volumes flexgroup &dash; FlexGroups.
svm	<a href="#">svm</a>	SVM containing the volume. Required on POST.
tiering	<a href="#">tiering</a>	
type	string	Type of the volume. rw &dash; read-write volume. dp &dash; data-protection volume. ls &dash; load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	<p>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.</p>
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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> <li>• readOnly: 1</li> </ul>



## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "application": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "autosize": {
    "mode": "grow"
  },
  "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
},
"comment": "string",
"create_time": "2018-06-04 19:00:00 UTC",
"efficiency": {
    "compaction": "inline",
    "compression": "inline",
    "cross_volume_dedupe": "inline",
    "dedupe": "inline"
},
"encryption": {
    "key_id": "string",
    "state": "encrypted",
    "status": {
        "code": "string",
        "message": "string"
    },
    "type": "none"
},
"files": {
    "used": 0
},
"flexcache_endpoint_type": "none",
"guarantee": {
    "type": "volume"
},
"language": "ar",
"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": "2017-01-25 11:20:13 UTC"
},
"movement": {
  "cutover_window": 30,
  "destination_aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "percent_complete": "string",
  "state": "replicating",
  "tiering_policy": "all"
},
"name": "vol_cs_dept",
"nas": {
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "path": "/user/my_volume",
  "security_style": "mixed",
  "unix_permissions": 493
},
"qos": {
  "policy": {
    "_links": {
```

```

    "self": {
      "href": "/api/resourcelink"
    }
  },
  "max_throughput_iops": 10000,
  "max_throughput_mbps": 500,
  "min_throughput_iops": 2000,
  "name": "performance",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
},
"quota": {
  "state": "corrupt"
},
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04 19:00:00 UTC",
  "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"
},
"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,
  "logical_space": {
    "available": 0,
    "used_by_afs": 0
  }
}
}

```

```

    },
    "metadata": 0,
    "over_provisioned": 0,
    "snapshot": {
      "used": 0
    },
    "used": 0
  },
  "state": "error",
  "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": "2017-01-25 11:20:13 UTC"
  },
  "style": "flexvol",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "policy": "all"
  },
  "type": "rw",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}

```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

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

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
918233	The target field cannot be specified for this operation.
2621707	No SVM was specified. Either "svm.name" or "svm.uuid" must be supplied.
918247	Specifying a value is not valid for a volume FlexClone creation.
918246	"volume.clone.parent_volume.name" or "volume.clone.parent_volume.uuid" must be provided.
2621706	The specified "svm.uuid" and "svm.name" do not refer to the same SVM.
918236	The specified "parent_volume.uuid" and "parent_volume.name" do not refer to the same volume.
918243	The specified Snapshot copy UUID is not correct for the specified Snapshot copy name.

Error Code	Description
918244	Invalid "volume.type" for clone volume.
918242	When creating a flexible volume, exactly one aggregate must be specified via either "aggregates.name" or "aggregates.uuid".
918241	The target style is an unsupported volume style for volume creation.
918240	The target style is an invalid volume style.
787140	One of "aggregates.uuid", "aggregates.name", or "style" must be provided.
787141	The specified "aggregates.name" and "aggregates.uuid" refer to different aggregates.
918252	"nas.path" is invalid.
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.

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	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

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 &dash; Volume automatically grows when the amount of used space is above the 'grow_threshold' value.  grow_shrink &dash; Volume grows or shrinks in response to the amount of space used.  off &dash; 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	<a href="#">_links</a>	
name	string	
uuid	string	

#### parent\_svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### parent\_volume

Name	Type	Description
_links	<a href="#">_links</a>	
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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

#### 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	<a href="#">snapshot_reference</a>	
parent_svm	<a href="#">parent_svm</a>	SVM, applies only to SVM-scoped objects.
parent_volume	<a href="#">parent_volume</a>	
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.

#### efficiency

Name	Type	Description
compaction	string	<p>The system can be enabled/disabled compaction.&lt;br&gt; inline &amp;dash;  Data will be compacted first and written to the volume.&lt;br&gt; none &amp;dash;  None&lt;br&gt; mixed &amp;dash;  Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.</p>
compression	string	<p>The system can be enabled/disabled compression.&lt;br&gt; inline &amp;dash;  Data will be compressed first and written to the volume.&lt;br&gt; background &amp;dash;  Data will be written to the volume and compressed later.&lt;br&gt; both &amp;dash;  Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run.&lt;br&gt; none &amp;dash;  None&lt;br&gt; mixed &amp;dash;  Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.</p>

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 &dash; Data will be cross volume deduped first and written to the volume.  background &dash; Data will be written to the volume and cross volume deduped later.  both &dash; 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 &dash; None  mixed &dash; 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 &dash; Data will be deduped first and written to the volume.  background &dash; Data will be written to the volume and deduped later.  both &dash; Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run.  none &dash; None  mixed &dash; Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.

status

Name	Type	Description
code	string	Encryption progress message code.
message	string	Encryption progress message.

encryption

Name	Type	Description
enabled	boolean	Encrypts an unencrypted volume. When set to 'true', a new key is generated and used to encrypt the given volume. The underlying SVM must be configured with the key manager.
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 &dash; The volume is completely encrypted.  encrypting &dash; Encryption operation is in progress.  partial &dash; 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 &dash; The volume is a plain-text one.
status	<a href="#">status</a>	
type	string	Volume encryption type.  none &dash; The volume is a plain-text one.  volume &dash; The volume is encrypted with volume key (NVE volume).  aggregate &dash; The volume is encrypted with aggregate key (NAE volume).

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 &dash; File system is inconsistent.  false &dash; File system in 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?

Name	Type	Description
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.
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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">iops</a>	The rate of I/O operations observed at the storage object.
latency	<a href="#">latency</a>	The round trip latency in microseconds observed at the storage object.

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 does not have the latest data.
throughput	<a href="#">throughput</a>	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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">destination_aggregate</a>	Aggregate
percent_complete	string	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. 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	<a href="#">_links</a>	
id	integer	
name	string	

nas

Name	Type	Description
export_policy	<a href="#">export_policy</a>	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. 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 &dash; Mixed-style security ntfs &dash; NTFS/Windows-style security unified &dash; Unified-style security, unified UNIX, NFS and CIFS permissions unix &dash; 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", "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
<a href="#">_links</a>	<a href="#">_links</a>	
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. Default is 15000 on AFF platforms and 10000 on all other platforms.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name, UUID and "min_throughput_iops" 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, UUID and "max_throughput_mbps" 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	<a href="#">policy</a>	When "min_throughput_iops", "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	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.

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&lt;num&gt;Y", "P&lt;num&gt;M", "P&lt;num&gt;D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT&lt;num&gt;H" and "PT&lt;num&gt;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	<p>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&lt;num&gt;Y", "P&lt;num&gt;M", "P&lt;num&gt;D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT&lt;num&gt;H" and "PT&lt;num&gt;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>

snaplock

Name	Type	Description
append_mode_enabled	boolean	<p>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.</p>

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	<a href="#">retention</a>	
type	string	The SnapLock type of the volume.   compliance &dash; 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 &dash; An administrator can delete a SnapLock Enterprise(SLE) volume.   non_snaplock &dash; Indicates the volume is non-snaplock.

#### snapshot\_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	<a href="#">_links</a>	
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.

Name	Type	Description
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.

Name	Type	Description
capacity_tier_footprint	integer	The space used by capacity tier for this volume in the aggregate, in bytes.
footprint	integer	Data and metadata used for this volume in the aggregate, in bytes.
logical_space	<a href="#">logical_space</a>	
metadata	integer	The space used by the total metadata in the volume, in bytes.
over_provisioned	integer	The amount of space not available for this volume in the aggregate, in bytes.
size	integer	Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes.
snapshot	<a href="#">snapshot</a>	
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 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.
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 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	<a href="#">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	<a href="#">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	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 does not have the latest data.
throughput_raw	<a href="#">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	The timestamp of the performance data.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

tiering

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 &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.   auto &dash; This policy allows tiering of both snapshot and active file system user data to the cloud store  none &dash; Volume blocks will not be tiered to the cloud store.  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.



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
_links	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	Aggregate hosting the volume. Required on POST.
application	<a href="#">application</a>	
autosize	<a href="#">autosize</a>	
clone	<a href="#">clone</a>	
comment	string	A comment for the volume. Valid in POST or PATCH.

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	<a href="#">efficiency</a>	
encryption	<a href="#">encryption</a>	
error_state	<a href="#">error_state</a>	
files	<a href="#">files</a>	
flexcache_endpoint_type	string	FlexCache endpoint type. none &dash; The volume is neither a FlexCache nor origin of any FlexCache. cache &dash; The volume is a FlexCache volume. origin &dash; The volume is origin of a FlexCache volume.
guarantee	<a href="#">guarantee</a>	
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	<a href="#">metric</a>	Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
movement	<a href="#">movement</a>	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	<a href="#">nas</a>	
qos	<a href="#">qos</a>	QoS information
quota	<a href="#">quota</a>	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. 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	<a href="#">snaplock</a>	
snapshot_policy	<a href="#">snapshot_policy</a>	This is a reference to the Snapshot copy policy.
space	<a href="#">space</a>	

Name	Type	Description
state	string	Volume state. A volume can only be brought online if it is offline. 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	<a href="#">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.
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. If "style" is specified, a volume of that type is created. That is, if style is "flexvol", a single aggregate must be specified. If style is "flexgroup", the system either uses the specified aggregates, or automatically provisions if no aggregates are specified. flexvol &dash; flexible volumes and FlexClone volumes flexgroup &dash; FlexGroups.
svm	<a href="#">svm</a>	SVM containing the volume. Required on POST.
tiering	<a href="#">tiering</a>	
type	string	Type of the volume. rw &dash; read-write volume. dp &dash; data-protection volume. ls &dash; 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	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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> <li>• readOnly: 1</li> </ul>

#### job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}
```

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.

### Related ONTAP commands

- `volume delete`
- `volume clone delete`

### Learn more

- [DOC /storage/volumes](#)

### Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

### Response

```
Status: 202, Accepted
```

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}



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 [DOC Requesting specific fields](#) to learn more.

- `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.footprint`

- `space.over_provisioned`
- `space.metadata`
- `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`

### Learn more

- [DOC /storage/volumes](#)

### Parameters

Name	Type	In	Required	Description
<code>uuid</code>	<code>string</code>	<code>path</code>	True	Unique identifier of the volume.
<code>fields</code>	<code>array[string]</code>	<code>query</code>	False	Specify the fields to return.

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	Aggregate hosting the volume. Required on POST.
application	<a href="#">application</a>	
autosize	<a href="#">autosize</a>	
clone	<a href="#">clone</a>	
comment	string	A comment for the volume. Valid in POST or PATCH.
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	<a href="#">efficiency</a>	
encryption	<a href="#">encryption</a>	
error_state	<a href="#">error_state</a>	
files	<a href="#">files</a>	
flexcache_endpoint_type	string	FlexCache endpoint type. none &dash; The volume is neither a FlexCache nor origin of any FlexCache. cache &dash; The volume is a FlexCache volume. origin &dash; The volume is origin of a FlexCache volume.
guarantee	<a href="#">guarantee</a>	

Name	Type	Description
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	<a href="#">metric</a>	Performance numbers, such as IOPS latency and throughput.
movement	<a href="#">movement</a>	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	<a href="#">nas</a>	
qos	<a href="#">qos</a>	QoS information
quota	<a href="#">quota</a>	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. 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	<a href="#">snaplock</a>	
snapshot_policy	<a href="#">snapshot_policy</a>	This is a reference to the Snapshot copy policy.

Name	Type	Description
space	space	
state	string	Volume state. A volume can only be brought online if it is offline. 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.
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. If "style" is specified, a volume of that type is created. That is, if style is "flexvol", a single aggregate must be specified. If style is "flexgroup", the system either uses the specified aggregates, or automatically provisions if no aggregates are specified. flexvol &dash; flexible volumes and FlexClone volumes flexgroup &dash; FlexGroups.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw &dash; read-write volume. dp &dash; data-protection volume. ls &dash; load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	<p>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.</p>
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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> <li>• readOnly: 1</li> </ul>

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "application": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "autosize": {
    "mode": "grow"
  },
  "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
},
"comment": "string",
"create_time": "2018-06-04 19:00:00 UTC",
"efficiency": {
    "compaction": "inline",
    "compression": "inline",
    "cross_volume_dedupe": "inline",
    "dedupe": "inline"
},
"encryption": {
    "key_id": "string",
    "state": "encrypted",
    "status": {
        "code": "string",
        "message": "string"
    },
    "type": "none"
},
"files": {
    "used": 0
},
"flexcache_endpoint_type": "none",
"guarantee": {
    "type": "volume"
},
"language": "ar",
"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": "2017-01-25 11:20:13 UTC"
},
"movement": {
  "cutover_window": 30,
  "destination_aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "percent_complete": "string",
  "state": "replicating",
  "tiering_policy": "all"
},
"name": "vol_cs_dept",
"nas": {
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "path": "/user/my_volume",
  "security_style": "mixed",
  "unix_permissions": 493
},
"qos": {
  "policy": {
    "_links": {

```

```
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "max_throughput_iops": 10000,
  "max_throughput_mbps": 500,
  "min_throughput_iops": 2000,
  "name": "performance",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
},
"quota": {
  "state": "corrupt"
},
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04 19:00:00 UTC",
  "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"
},
"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,
  "logical_space": {
    "available": 0,
    "used_by_afs": 0
  }
}
```

```

    },
    "metadata": 0,
    "over_provisioned": 0,
    "snapshot": {
      "used": 0
    },
    "used": 0
  },
  "state": "error",
  "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": "2017-01-25 11:20:13 UTC"
  },
  "style": "flexvol",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "policy": "all"
  },
  "type": "rw",
  "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	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

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 &dash; Volume automatically grows when the amount of used space is above the 'grow_threshold' value.  grow_shrink &dash; Volume grows or shrinks in response to the amount of space used.  off &dash; 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	<a href="#">_links</a>	
name	string	
uuid	string	

#### parent\_svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### parent\_volume

Name	Type	Description
_links	<a href="#">_links</a>	
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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

#### 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	<a href="#">snapshot_reference</a>	
parent_svm	<a href="#">parent_svm</a>	SVM, applies only to SVM-scoped objects.
parent_volume	<a href="#">parent_volume</a>	
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.

#### efficiency



Name	Type	Description
compaction	string	The system can be enabled/disabled compaction.  inline &dash; Data will be compacted first and written to the volume.  none &dash; None  mixed &dash; 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 &dash; Data will be compressed first and written to the volume.  background &dash; Data will be written to the volume and compressed later.  both &dash; Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run.  none &dash; None  mixed &dash; 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 &dash; Data will be cross volume deduped first and written to the volume.  background &dash; Data will be written to the volume and cross volume deduped later.  both &dash; 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 &dash; None  mixed &dash; 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 &dash; Data will be deduped first and written to the volume.  background &dash; Data will be written to the volume and deduped later.  both &dash; Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run.  none &dash; None  mixed &dash; Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.

status

Name	Type	Description
code	string	Encryption progress message code.
message	string	Encryption progress message.

encryption

Name	Type	Description
enabled	boolean	Encrypts an unencrypted volume. When set to 'true', a new key is generated and used to encrypt the given volume. The underlying SVM must be configured with the key manager.
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 &dash; The volume is completely encrypted.  encrypting &dash; Encryption operation is in progress.  partial &dash; 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 &dash; The volume is a plain-text one.
status	<a href="#">status</a>	
type	string	Volume encryption type.  none &dash; The volume is a plain-text one.  volume &dash; The volume is encrypted with volume key (NVE volume).  aggregate &dash; The volume is encrypted with aggregate key (NAE volume).

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 &dash; File system is inconsistent.  false &dash; File system in 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?

Name	Type	Description
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.
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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">iops</a>	The rate of I/O operations observed at the storage object.
latency	<a href="#">latency</a>	The round trip latency in microseconds observed at the storage object.

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 does not have the latest data.
throughput	<a href="#">throughput</a>	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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">destination_aggregate</a>	Aggregate
percent_complete	string	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. 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	<a href="#">_links</a>	
id	integer	
name	string	

nas

Name	Type	Description
export_policy	<a href="#">export_policy</a>	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. 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 &dash; Mixed-style security ntfs &dash; NTFS/Windows-style security unified &dash; Unified-style security, unified UNIX, NFS and CIFS permissions unix &dash; 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", "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	<a href="#">_links</a>	
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. Default is 15000 on AFF platforms and 10000 on all other platforms.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name, UUID and "min_throughput_iops" 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, UUID and "max_throughput_mbps" 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	<a href="#">policy</a>	When "min_throughput_iops", "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	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.

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&lt;num&gt;Y", "P&lt;num&gt;M", "P&lt;num&gt;D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT&lt;num&gt;H" and "PT&lt;num&gt;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	<p>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&lt;num&gt;Y", "P&lt;num&gt;M", "P&lt;num&gt;D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT&lt;num&gt;H" and "PT&lt;num&gt;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>

snaplock

Name	Type	Description
append_mode_enabled	boolean	<p>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.</p>

Name	Type	Description
autocommit_period	string	<p>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&lt;num&gt;Y", "P&lt;num&gt;M", "P&lt;num&gt;D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT&lt;num&gt;H" and "PT&lt;num&gt;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".</p>
compliance_clock_time	string	<p>This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.</p>
expiry_time	string	<p>Expiry time of the volume.</p>
is_audit_log	boolean	<p>Indicates if this volume has been configured as SnapLock audit log volume for the SVM .</p>
litigation_count	integer	<p>Litigation count indicates the number of active legal-holds on the volume.</p>

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	<a href="#">retention</a>	
type	string	The SnapLock type of the volume.   compliance &dash; 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 &dash; An administrator can delete a SnapLock Enterprise(SLE) volume.   non_snaplock &dash; Indicates the volume is non-snaplock.

#### snapshot\_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	<a href="#">_links</a>	
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.



Name	Type	Description
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.

Name	Type	Description
capacity_tier_footprint	integer	The space used by capacity tier for this volume in the aggregate, in bytes.
footprint	integer	Data and metadata used for this volume in the aggregate, in bytes.
logical_space	<a href="#">logical_space</a>	
metadata	integer	The space used by the total metadata in the volume, in bytes.
over_provisioned	integer	The amount of space not available for this volume in the aggregate, in bytes.
size	integer	Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes.
snapshot	<a href="#">snapshot</a>	
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 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.
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 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	<a href="#">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	<a href="#">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	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 does not have the latest data.
throughput_raw	<a href="#">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	The timestamp of the performance data.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

tiering

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 &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.   auto &dash; This policy allows tiering of both snapshot and active file system user data to the cloud store  none &dash; Volume blocks will not be tiered to the cloud store.  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.

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.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

### Update volume attributes

PATCH /storage/volumes/{uuid}

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.

#### Related ONTAP commands

- `volume modify`
- `volume clone modify`
- `volume efficiency modify`
- `volume quota on`
- `volume quota off`

#### Learn more

- [DOC /storage/volumes](#)

#### Parameters

Name	Type	In	Required	Description
<code>uuid</code>	string	path	True	Unique identifier of the volume.
<code>restore_to.snapshot.uuid</code>	string	query	False	UUID of the Snapshot copy to restore volume to the point in time the Snapshot copy was taken.
<code>restore_to.snapshot.name</code>	string	query	False	Name of the Snapshot copy to restore volume to the point in time the Snapshot copy was taken.

Name	Type	In	Required	Description
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"> <li>• use_existing_resources - Increases or decreases the size of the FlexGroup by increasing or decreasing the size of the current FlexGroup resources</li> <li>• add_new_resources - Increases the size of the FlexGroup by adding new resources</li> <li>• Default value: 1</li> <li>• enum: ["use_existing_resources", "add_new_resources"]</li> </ul>
movement.destination_aggregate.name	string	query	False	Name of the aggregates to which the specified volume can be moved.
movement.destination_aggregate.uuid	string	query	False	UUID of the aggregates to which the specified volume can be moved.
validate_only	boolean	query	False	Validate the operation and its parameters, without actually performing the operation.



## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	Aggregate hosting the volume. Required on POST.
application	<a href="#">application</a>	
autosize	<a href="#">autosize</a>	
clone	<a href="#">clone</a>	
comment	string	A comment for the volume. Valid in POST or PATCH.
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	<a href="#">efficiency</a>	
encryption	<a href="#">encryption</a>	
error_state	<a href="#">error_state</a>	
files	<a href="#">files</a>	
flexcache_endpoint_type	string	FlexCache endpoint type.  none &dash; The volume is neither a FlexCache nor origin of any FlexCache.  cache &dash; The volume is a FlexCache volume.  origin &dash; The volume is origin of a FlexCache volume.
guarantee	<a href="#">guarantee</a>	

Name	Type	Description
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	<a href="#">metric</a>	Performance numbers, such as IOPS latency and throughput.
movement	<a href="#">movement</a>	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	<a href="#">nas</a>	
qos	<a href="#">qos</a>	QoS information
quota	<a href="#">quota</a>	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. 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	<a href="#">snaplock</a>	
snapshot_policy	<a href="#">snapshot_policy</a>	This is a reference to the Snapshot copy policy.

Name	Type	Description
space	space	
state	string	Volume state. A volume can only be brought online if it is offline. 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.
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. If "style" is specified, a volume of that type is created. That is, if style is "flexvol", a single aggregate must be specified. If style is "flexgroup", the system either uses the specified aggregates, or automatically provisions if no aggregates are specified. flexvol &dash; flexible volumes and FlexClone volumes flexgroup &dash; FlexGroups.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw &dash; read-write volume. dp &dash; data-protection volume. ls &dash; load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	<p>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.</p>
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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> <li>• readOnly: 1</li> </ul>

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "application": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "autosize": {
    "mode": "grow"
  },
  "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
},
"comment": "string",
"create_time": "2018-06-04 19:00:00 UTC",
"efficiency": {
  "compaction": "inline",
  "compression": "inline",
  "cross_volume_dedupe": "inline",
  "dedupe": "inline"
},
"encryption": {
  "key_id": "string",
  "state": "encrypted",
  "status": {
    "code": "string",
    "message": "string"
  },
  "type": "none"
},
"files": {
  "used": 0
},
"flexcache_endpoint_type": "none",
"guarantee": {
  "type": "volume"
},
"language": "ar",
"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": "2017-01-25 11:20:13 UTC"
  },
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": "string",
    "state": "replicating",
    "tiering_policy": "all"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 100,
      "name": "default"
    },
    "path": "/user/my_volume",
    "security_style": "mixed",
    "unix_permissions": 493
  },
  "qos": {
    "policy": {
      "_links": {

```

```

    "self": {
      "href": "/api/resourcelink"
    }
  },
  "max_throughput_iops": 10000,
  "max_throughput_mbps": 500,
  "min_throughput_iops": 2000,
  "name": "performance",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
},
"quota": {
  "state": "corrupt"
},
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04 19:00:00 UTC",
  "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"
},
"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,
  "logical_space": {
    "available": 0,
    "used_by_afs": 0
  }
}
}

```



```
    },
    "metadata": 0,
    "over_provisioned": 0,
    "snapshot": {
      "used": 0
    },
    "used": 0
  },
  "state": "error",
  "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": "2017-01-25 11:20:13 UTC"
  },
  "style": "flexvol",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "policy": "all"
  },
  "type": "rw",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

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

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
13109198	Resizing by adding new resources is only supported for FlexGroups.
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.
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.
787141	The specified "aggregates.name" and "aggregates.uuid" refer to different aggregates.

Error Code	Description
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.
918265	Volume is on the same aggregate.
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.

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	<a href="#">href</a>	

aggregates

Aggregate

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

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 &dash; Volume automatically grows when the amount of used space is above the 'grow_threshold' value.  grow_shrink &dash; Volume grows or shrinks in response to the amount of space used.  off &dash; 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	<a href="#">_links</a>	
name	string	
uuid	string	

#### parent\_svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

#### parent\_volume

Name	Type	Description
_links	<a href="#">_links</a>	
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"> <li>example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> </ul>

#### 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	<a href="#">snapshot_reference</a>	
parent_svm	<a href="#">parent_svm</a>	SVM, applies only to SVM-scoped objects.
parent_volume	<a href="#">parent_volume</a>	
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.

#### efficiency

Name	Type	Description
compaction	string	The system can be enabled/disabled compaction.  inline &dash; Data will be compacted first and written to the volume.  none &dash; None  mixed &dash; 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 &dash; Data will be compressed first and written to the volume.  background &dash; Data will be written to the volume and compressed later.  both &dash; Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run.  none &dash; None  mixed &dash; 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 &dash; Data will be cross volume deduped first and written to the volume.  background &dash; Data will be written to the volume and cross volume deduped later.  both &dash; 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 &dash; None  mixed &dash; 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 &dash; Data will be deduped first and written to the volume.  background &dash; Data will be written to the volume and deduped later.  both &dash; Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run.  none &dash; None  mixed &dash; Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.

status

Name	Type	Description
code	string	Encryption progress message code.
message	string	Encryption progress message.

encryption

Name	Type	Description
enabled	boolean	Encrypts an unencrypted volume. When set to 'true', a new key is generated and used to encrypt the given volume. The underlying SVM must be configured with the key manager.
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 &dash; The volume is completely encrypted.  encrypting &dash; Encryption operation is in progress.  partial &dash; 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 &dash; The volume is a plain-text one.
status	<a href="#">status</a>	
type	string	Volume encryption type.  none &dash; The volume is a plain-text one.  volume &dash; The volume is encrypted with volume key (NVE volume).  aggregate &dash; The volume is encrypted with aggregate key (NAE volume).

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 &dash; File system is inconsistent.  false &dash; File system in 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?

Name	Type	Description
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.
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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">iops</a>	The rate of I/O operations observed at the storage object.
latency	<a href="#">latency</a>	The round trip latency in microseconds observed at the storage object.

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 does not have the latest data.
throughput	<a href="#">throughput</a>	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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">destination_aggregate</a>	Aggregate
percent_complete	string	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. 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	<a href="#">_links</a>	
id	integer	
name	string	

nas

Name	Type	Description
export_policy	<a href="#">export_policy</a>	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. 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 &dash; Mixed-style security ntfs &dash; NTFS/Windows-style security unified &dash; Unified-style security, unified UNIX, NFS and CIFS permissions unix &dash; 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", "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
<a href="#">_links</a>	<a href="#">_links</a>	
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. Default is 15000 on AFF platforms and 10000 on all other platforms.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name, UUID and "min_throughput_iops" 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, UUID and "max_throughput_mbps" 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	<a href="#">policy</a>	When "min_throughput_iops", "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	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.

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&lt;num&gt;Y", "P&lt;num&gt;M", "P&lt;num&gt;D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT&lt;num&gt;H" and "PT&lt;num&gt;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	<p>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&lt;num&gt;Y", "P&lt;num&gt;M", "P&lt;num&gt;D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT&lt;num&gt;H" and "PT&lt;num&gt;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>

snaplock

Name	Type	Description
append_mode_enabled	boolean	<p>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.</p>

Name	Type	Description
autocommit_period	string	<p>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&lt;num&gt;Y", "P&lt;num&gt;M", "P&lt;num&gt;D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT&lt;num&gt;H" and "PT&lt;num&gt;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".</p>
compliance_clock_time	string	<p>This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.</p>
expiry_time	string	<p>Expiry time of the volume.</p>
is_audit_log	boolean	<p>Indicates if this volume has been configured as SnapLock audit log volume for the SVM .</p>
litigation_count	integer	<p>Litigation count indicates the number of active legal-holds on the volume.</p>

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	<a href="#">retention</a>	
type	string	The SnapLock type of the volume.   compliance &dash; 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 &dash; An administrator can delete a SnapLock Enterprise(SLE) volume.   non_snaplock &dash; Indicates the volume is non-snaplock.

#### snapshot\_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	<a href="#">_links</a>	
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.

Name	Type	Description
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.



Name	Type	Description
capacity_tier_footprint	integer	The space used by capacity tier for this volume in the aggregate, in bytes.
footprint	integer	Data and metadata used for this volume in the aggregate, in bytes.
logical_space	<a href="#">logical_space</a>	
metadata	integer	The space used by the total metadata in the volume, in bytes.
over_provisioned	integer	The amount of space not available for this volume in the aggregate, in bytes.
size	integer	Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes.
snapshot	<a href="#">snapshot</a>	
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 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.
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 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 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	<p>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".</p> <p>"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.</p> <p>"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.</p>
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

SVM containing the volume. Required on POST.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

tiering

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 &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.   auto &dash; This policy allows tiering of both snapshot and active file system user data to the cloud store  none &dash; Volume blocks will not be tiered to the cloud store.  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.

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
_links	<a href="#">_links</a>	
aggregates	array[ <a href="#">aggregates</a> ]	Aggregate hosting the volume. Required on POST.
application	<a href="#">application</a>	
autosize	<a href="#">autosize</a>	
clone	<a href="#">clone</a>	
comment	string	A comment for the volume. Valid in POST or PATCH.

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	<a href="#">efficiency</a>	
encryption	<a href="#">encryption</a>	
error_state	<a href="#">error_state</a>	
files	<a href="#">files</a>	
flexcache_endpoint_type	string	FlexCache endpoint type. none &dash; The volume is neither a FlexCache nor origin of any FlexCache. cache &dash; The volume is a FlexCache volume. origin &dash; The volume is origin of a FlexCache volume.
guarantee	<a href="#">guarantee</a>	
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	<a href="#">metric</a>	Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
movement	<a href="#">movement</a>	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	<a href="#">nas</a>	
qos	<a href="#">qos</a>	QoS information
quota	<a href="#">quota</a>	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. 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	<a href="#">snaplock</a>	
snapshot_policy	<a href="#">snapshot_policy</a>	This is a reference to the Snapshot copy policy.
space	<a href="#">space</a>	

Name	Type	Description
state	string	Volume state. A volume can only be brought online if it is offline. 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	<a href="#">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.
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. If "style" is specified, a volume of that type is created. That is, if style is "flexvol", a single aggregate must be specified. If style is "flexgroup", the system either uses the specified aggregates, or automatically provisions if no aggregates are specified. flexvol &dash; flexible volumes and FlexClone volumes flexgroup &dash; FlexGroups.
svm	<a href="#">svm</a>	SVM containing the volume. Required on POST.
tiering	<a href="#">tiering</a>	
type	string	Type of the volume. rw &dash; read-write volume. dp &dash; data-protection volume. ls &dash; 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"> <li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li> <li>• readOnly: 1</li> </ul>

#### job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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/{uuid}/metrics

Retrieves historical performance metrics for a volume.

### Parameters

Name	Type	In	Required	Description
status	string	query	False	Filter by status
timestamp	string	query	False	Filter by timestamp
duration	string	query	False	Filter by duration
iops.total	integer	query	False	Filter by iops.total
iops.other	integer	query	False	Filter by iops.other
iops.read	integer	query	False	Filter by iops.read
iops.write	integer	query	False	Filter by iops.write
latency.total	integer	query	False	Filter by latency.total
latency.other	integer	query	False	Filter by latency.other
latency.read	integer	query	False	Filter by latency.read
latency.write	integer	query	False	Filter by latency.write

Name	Type	In	Required	Description
throughput.total	integer	query	False	Filter by throughput.total
throughput.other	integer	query	False	Filter by throughput.other
throughput.read	integer	query	False	Filter by throughput.read
throughput.write	integer	query	False	Filter by throughput.write
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"> <li>• Default value: 1</li> </ul>
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
<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"> <li>• Default value: 1</li> </ul>	uuid	string	path	True

Name	Type	In	Required	Description
Unique identifier of the volume.	interval	string	query	False

## Response

Status: 200, Ok

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">records</a> ]	

## 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": "2017-01-25 11:20:13 UTC"
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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
<a href="#">_links</a>	<a href="#">_links</a>	



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 does 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[ <a href="#">error_arguments</a> ]	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 point in 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.

&ndash; POST /api/storage/volumes/{volume.uuid}/snapshots

&ndash; GET /api/storage/volumes/{volume.uuid}/snapshots

&ndash; GET /api/storage/volumes/{volume.uuid}/snapshots/{uuid}

&ndash; PATCH /api/storage/volumes/{volume.uuid}/snapshots/{uuid}

&ndash; 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

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	False	Volume
volume.name	string	query	False	Filter by volume.name
volume.uuid	string	query	False	Filter by volume.uuid
state	string	query	False	Filter by state
name	string	query	False	Filter by name
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
create_time	string	query	False	Filter by create_time
snaplock_expiry_time	string	query	False	Filter by snaplock_expiry_time
expiry_time	string	query	False	Filter by expiry_time
uuid	string	query	False	Filter by uuid
comment	string	query	False	Filter by comment
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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">snapshot</a> ]	

## Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "comment": "string",
    "create_time": "2019-02-04 19:00:00 UTC",
    "expiry_time": "2019-02-04 19:00:00 UTC",
    "name": "this_snapshot",
    "snaplock_expiry_time": "2019-02-04 19:00:00 UTC",
    "state": "valid",
    "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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li></ul>

snapshot

The Snapshot copy object represents a point in time Snapshot copy of a volume.

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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.
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.



Name	Type	Description
svm	svm	SVM, applies only to SVM-scoped objects.
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

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

## Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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.
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	SVM, applies only to SVM-scoped objects.
uuid	string	The UUID of the Snapshot copy in the volume that uniquely identifies the Snapshot copy in that volume.
volume	volume	

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "create_time": "2019-02-04 19:00:00 UTC",
  "expiry_time": "2019-02-04 19:00:00 UTC",
  "name": "this_snapshot",
  "snaplock_expiry_time": "2019-02-04 19:00:00 UTC",
  "state": "valid",
  "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"
  }
}
```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li></ul>

snapshot

The Snapshot copy object represents a point in time Snapshot copy of a volume.

Name	Type	Description
_links	<a href="#">_links</a>	

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.
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	SVM, applies only to SVM-scoped objects.

Name	Type	Description
uuid	string	The UUID of the Snapshot copy in the volume that uniquely identifies the Snapshot copy in that volume.
volume	<a href="#">volume</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}
```

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	
uuid	string	path	True	

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## Example response

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

## Error

Status: Default, Error

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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
<a href="#">_links</a>	<a href="#">_links</a>	
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.

Name	Type	Description
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	SVM, applies only to SVM-scoped objects.
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": "2019-02-04 19:00:00 UTC",
  "expiry_time": "2019-02-04 19:00:00 UTC",
  "name": "this_snapshot",
  "snaplock_expiry_time": "2019-02-04 19:00:00 UTC",
  "state": "valid",
  "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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li></ul>

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument



error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}

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

#### Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
comment	string	A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.

Name	Type	Description
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.
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	SVM, applies only to SVM-scoped objects.
uuid	string	The UUID of the Snapshot copy in the volume that uniquely identifies the Snapshot copy in that volume.
volume	volume	

## Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "create_time": "2019-02-04 19:00:00 UTC",
  "expiry_time": "2019-02-04 19:00:00 UTC",
  "name": "this_snapshot",
  "snaplock_expiry_time": "2019-02-04 19:00:00 UTC",
  "state": "valid",
  "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"
  }
}
```

## Response

Status: 202, Accepted

Name	Type	Description
job	<a href="#">job_link</a>	

## 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	<a href="#">href</a>	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	<a href="#">_links</a>	
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"><li>• example: 028baa66-41bd-11e9-81d5-00a0986138f7</li></ul>

snapshot

The Snapshot copy object represents a point in time Snapshot copy of a volume.

Name	Type	Description
_links	<a href="#">_links</a>	

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.
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	SVM, applies only to SVM-scoped objects.

Name	Type	Description
uuid	string	The UUID of the Snapshot copy in the volume that uniquely identifies the Snapshot copy in that volume.
volume	<a href="#">volume</a>	

job\_link

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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 NetApp's *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. GET calls can also be used 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" OR
curl -X GET "https://<mgmt-ip>/api/support/autosupport?fields=*" -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 any 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=*&issues"
-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

Retrieves the AutoSupport configuration of the cluster and if requested, returns any 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 [DOC 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
fields	array[string]	query	False	Specify the fields to return.

### 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.
from	string	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. <ul style="list-style-type: none"> <li>example: <a href="mailto:postmaster@example.com">postmaster@example.com</a></li> <li>format: email</li> </ul>
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[ <a href="#">autosupport_issues</a> ]	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.

<b>Name</b>	<b>Type</b>	<b>Description</b>
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	Corrective action message. The corrective action might contain commands which needs to be executed on the ONTAP CLI. <ul style="list-style-type: none"><li>• example: Check the hostname of the SMTP server</li><li>• readOnly: 1</li></ul>

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	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

autosupport\_issues

Name	Type	Description
corrective_action	<a href="#">autosupport_connectivity_corrective_action</a>	
issue	<a href="#">autosupport_connectivity_issue</a>	
node	<a href="#">node</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update AutoSupport configuration for an entire cluster

PATCH /support/autosupport

Updates the AutoSupport configuration for the entire cluster.

### Related ONTAP commands

- `system node autosupport modify`

### Learn more

- [DOC /support/autosupport](#)

### 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	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. <ul style="list-style-type: none"> <li>example: <a href="mailto:postmaster@example.com">postmaster@example.com</a></li> <li>format: email</li> </ul>
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[ <a href="#">autosupport_issues</a> ]	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.

<b>Name</b>	<b>Type</b>	<b>Description</b>
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": {
      "_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"
}
```

## 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

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	Corrective action message. The corrective action might contain commands which needs to be executed on the ONTAP CLI. <ul style="list-style-type: none"><li>• example: Check the hostname of the SMTP server</li><li>• readOnly: 1</li></ul>

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	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

autosupport\_issues

Name	Type	Description
corrective_action	<a href="#">autosupport_connectivity_corrective_action</a>	
issue	<a href="#">autosupport_connectivity_issue</a>	
node	<a href="#">node</a>	

#### 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	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. <ul style="list-style-type: none"> <li>example: <a href="mailto:postmaster@example.com">postmaster@example.com</a></li> <li>format: email</li> </ul>
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[ <a href="#">autosupport_issues</a> ]	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.

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

This API is used to invoke and retrieve AutoSupport messages from the nodes in the cluster.

This API supports POST and GET calls. A POST call is used to invoke AutoSupport and a GET call is used 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` param 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` param 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\": \"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\_successful' 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. Similar to this, filters can be added 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

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
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	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.

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">autosupport_message</a> ]	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": "2019-03-25 17:30:04 -0400",
    "index": 9,
    "message": "invoked_test_autosupport_rest",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "state": "sent_successful",
    "subject": "WEEKLY_LOG",
    "type": "test",
    "uri": "http://1.2.3.4/delivery_uri"
  }
}
```

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

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	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

autosupport\_message

Name	Type	Description
destination	string	Destination for the AutoSupport
error	<a href="#">error</a>	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
message	string	Message included in the AutoSupport subject
node	<a href="#">node</a>	
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

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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

Generates 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 to identify the generated AutoSupport message.

#### Default property values

If not specified in POST, the following default property values are assigned:

- `type` - *all*
- `node.name` or `node.uuid` - Not specifying any of these properties invokes the 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.

#### Request Body

Name	Type	Description
<code>destination</code>	string	Destination for the AutoSupport
<code>error</code>	<a href="#">error</a>	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

Name	Type	Description
message	string	Message included in the AutoSupport subject
node	<a href="#">node</a>	
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": "2019-03-25 17:30:04 -0400",
  "index": 9,
  "message": "invoked_test_autosupport_rest",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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
<code>_links</code>	<code>_links</code>	
<code>num_records</code>	integer	Number of records
<code>records</code>	array[ <code>records</code> ]	

## Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 3,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "index": 9,
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "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

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

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

Name	Type	Description
self	<a href="#">href</a>	

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

autosupport\_message

Name	Type	Description
destination	string	Destination for the AutoSupport
error	<a href="#">error</a>	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



Name	Type	Description
node	<a href="#">node</a>	
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

#### records

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
index	integer	Sequence number of the generated AutoSupport
node	<a href="#">node</a>	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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": "http://10.224.65.198/new_backups",
  "username": "new_me",
  "password": "new_pass"
}

# 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

Retrieves the cluster configuration backup information.

### Learn more

- [DOC /support/configuration-backup](#)

### Parameters

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

### 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	

### 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[ <a href="#">error_arguments</a> ]	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

Updates the cluster configuration backup information.

### Learn more

- [DOC /support/configuration-backup](#)

### Request Body

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	

## 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

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

#### 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	

#### error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

#### error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
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 e-mail destination

The following example configures EMS to send a support e-mail 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 EMS configuration

GET /support/ems

Retrieves the EMS configuration.

## Related ONTAP commands

- `event config show`

## Learn more

- [DOC /support/ems](#)

## Parameters

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

## Response

Status: 200, Ok

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>mail_from</code>	string	Mail from
<code>mail_server</code>	string	Mail server (SMTP)
<code>proxy_password</code>	string	Password for HTTP/HTTPS proxy
<code>proxy_url</code>	string	HTTP/HTTPS proxy URL
<code>proxy_user</code>	string	User name for HTTP/HTTPS proxy

## 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"
}
```

## 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	<a href="#">href</a>	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Update EMS configuration

PATCH `/support/ems`

Updates the EMS configuration.

### Related ONTAP commands

- `event config modify`

### Learn more

- [DOC /support/ems](#)

## Request Body

Name	Type	Description
_links	<a href="#">_links</a>	
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

## Example request

```
{
  "_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"
}
```

## 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'

Name	Type	Description
error	<a href="#">error</a>	

### 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	<a href="#">href</a>	

ems\_config

Name	Type	Description
_links	<a href="#">_links</a>	
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

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
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

#### E-mail

The 'email' type allows you to define a mail box where information about an observed event is sent via SMTP. The address should be populated in the destination field and be a valid format.

For example: [administrator@mycompany.com](mailto:administrator@mycompany.com)

#### Syslog

The 'syslog' type allows you to specify a remote syslog server that can receive information about an observed event. The hostname or IP address should be populated 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. The URL should be populated 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 if mutual authentication is desired.  
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

Retrieves a collection of event destinations.

#### Related ONTAP commands

- event notification destination show
- event notification show

#### Learn more

- [DOC /support/ems/destinations](#)

#### Parameters

Name	Type	In	Required	Description
certificate.serial_number	string	query	False	Filter by certificate.serial_number
certificate.ca	string	query	False	Filter by certificate.ca
destination	string	query	False	Filter by destination
filters.name	string	query	False	Filter by filters.name
name	string	query	False	Filter by name

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_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.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">records</a> ]	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

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	<a href="#">_links</a>	
name	string	

records

Name	Type	Description
_links	<a href="#">_links</a>	
certificate	<a href="#">certificate</a>	Certificate information is valid for the "rest_api" type.
destination	string	Event destination
filters	array[ <a href="#">filters</a> ]	

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[ <a href="#">error_arguments</a> ]	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

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 field 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`

### Learn more

- [DOC /support/ems/destinations](#)

### Request Body

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>certificate</code>	<code>certificate</code>	Certificate information is valid for the "rest_api" type.
<code>destination</code>	string	Event destination
<code>filters</code>	array[ <code>filters</code> ]	
<code>name</code>	string	Destination name. Valid in POST.
<code>type</code>	string	Type of destination. Valid in POST.



## Example request

```
{
  "_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"
}
```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">records</a> ]	

## 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

## 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 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

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	<a href="#">href</a>	

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	<a href="#">_links</a>	
name	string	

ems\_destination

Name	Type	Description
_links	<a href="#">_links</a>	
certificate	<a href="#">certificate</a>	Certificate information is valid for the "rest_api" type.
destination	string	Event destination
filters	array[ <a href="#">filters</a> ]	
name	string	Destination name. Valid in POST.
type	string	Type of destination. Valid in POST.

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

records

Name	Type	Description
_links	<a href="#">_links</a>	
certificate	<a href="#">certificate</a>	Certificate information is valid for the "rest_api" type.
destination	string	Event destination
filters	array[ <a href="#">filters</a> ]	
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[ <a href="#">error_arguments</a> ]	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 can be modified once a destination is created. For example, a destination's type cannot be changed but the underlying details of the type can be modified.

See the documentation for </support/ems/destinations> for details on the various fields 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 e-mail 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}

Deletes an event destination.

#### Related ONTAP commands

- event notification destination delete
- event notification delete

#### Learn more

- [DOC /support/ems/destinations/{name}](#)

#### Parameters

Name	Type	In	Required	Description
name	string	path	True	

#### 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	<a href="#">error</a>	

## 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[ <a href="#">error_arguments</a> ]	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}

Retrieves event destinations.

### Related ONTAP commands

- `event notification destination show`
- `event notification show`

### Learn more

- [DOC /support/ems/destinations/{name}](#)

### 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
<code>_links</code>	<code>_links</code>	
<code>certificate</code>	<code>certificate</code>	Certificate information is valid for the "rest_api" type.
<code>destination</code>	string	Event destination
<code>filters</code>	array[ <code>filters</code> ]	
<code>name</code>	string	Destination name. Valid in POST.
<code>type</code>	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	<a href="#">href</a>	

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	<a href="#">_links</a>	
name	string	

error\_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[ <a href="#">error_arguments</a> ]	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}

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`

#### Learn more

- [DOC /support/ems/destinations/{name}](#)

#### Parameters

Name	Type	In	Required	Description
name	string	path	True	Destination name

#### Request Body

Name	Type	Description
<code>_links</code>	<a href="#">_links</a>	
<code>certificate</code>	<a href="#">certificate</a>	Certificate information is valid for the "rest_api" type.
<code>destination</code>	string	Event destination
<code>filters</code>	array[ <a href="#">filters</a> ]	
<code>name</code>	string	Destination name. Valid in POST.
<code>type</code>	string	Type of destination. Valid in POST.

## Example request

```
{
  "_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"
}
```

## Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
certificate	<a href="#">certificate</a>	Certificate information is valid for the "rest_api" type.
destination	string	Event destination
filters	array[ <a href="#">filters</a> ]	
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

## 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



Error Code	Description
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
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

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	<a href="#">href</a>	

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	<a href="#">_links</a>	
name	string	

ems\_destination

Name	Type	Description
_links	<a href="#">_links</a>	
certificate	<a href="#">certificate</a>	Certificate information is valid for the "rest_api" type.
destination	string	Event destination
filters	array[ <a href="#">filters</a> ]	
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[ <a href="#">error_arguments</a> ]	Message arguments
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

Retrieves a collection of observed events.

## Related ONTAP commands

- `event log show`

## Learn more

- [DOC /support/ems/events](#)

## Parameters

Name	Type	In	Required	Description
source	string	query	False	Filter by source
log_message	string	query	False	Filter by log_message
message.name	string	query	False	Filter by message.name
message.severity	string	query	False	Filter by message.severity
index	integer	query	False	Filter by index
time	string	query	False	Filter by time
parameters.name	string	query	False	Filter by parameters.name
parameters.value	string	query	False	Filter by parameters.value
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.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.

Name	Type	In	Required	Description
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">records</a> ]	

**Example response**

A large, empty rectangular box with a thin, dashed border, occupying most of the page. It is intended for an 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

message

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Message name of the event. Returned by default.
severity	string	Severity of the event. Returned by default.

node

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	
uuid	string	

parameters

Name	Type	Description
name	string	Name of parameter
value	string	Value of parameter

records

Name	Type	Description
_links	<a href="#">_links</a>	
index	integer	Index of the event. Returned by default.
log_message	string	A formatted text string populated with parameter details. Returned by default.
message	<a href="#">message</a>	
node	<a href="#">node</a>	
parameters	array[ <a href="#">parameters</a> ]	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[ <a href="#">error_arguments</a> ]	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 an event is processed by EMS, 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

Retrieves a collection of event filters.

### Related ONTAP commands

- `event filter show`

### Learn more

- [DOC /support/ems/filters](#)

### Parameters

Name	Type	In	Required	Description
name	string	query	False	Filter by name
rules.message_criteria.severities	string	query	False	Filter by rules.message_criteria.severities



Name	Type	In	Required	Description
rules.message_criteria.name_pattern	string	query	False	Filter by rules.message_criteria.name_pattern
rules.message_criteria.snmp_trap_types	string	query	False	Filter by rules.message_criteria.snmp_trap_types
rules.type	string	query	False	Filter by rules.type
rules.index	integer	query	False	Filter by rules.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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">records</a> ]	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

\_links

Name	Type	Description
related	<a href="#">href</a>	

message\_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	

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	<a href="#">message_criteria</a>	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

#### records

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
name	string	Filter name
rules	array[ <a href="#">rules</a> ]	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[ <a href="#">error_arguments</a> ]	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

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`

### Learn more

- [DOC /support/ems/filters](#)

### Request Body

Name	Type	Description
<code>_links</code>	<code><a href="#">_links</a></code>	
<code>name</code>	string	Filter name
<code>rules</code>	array[ <code><a href="#">rules</a></code> ]	Array of event filter rules on which to match.

## Example request

```
{
  "_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"
  }
}
```

## Response

Status: 201, Created

Name	Type	Description
<code>_links</code>	<code><a href="#">_links</a></code>	
<code>num_records</code>	integer	Number of records
<code>records</code>	array[ <code><a href="#">records</a></code> ]	



## 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

## 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 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

Name	Type	Description
error	<a href="#">error</a>	

## 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	<a href="#">href</a>	

\_links

Name	Type	Description
related	<a href="#">href</a>	

message\_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
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	<a href="#">message_criteria</a>	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

#### ems\_filter

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Filter name
rules	array[ <a href="#">rules</a> ]	Array of event filter rules on which to match.

#### \_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

#### records

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Filter name
rules	array[ <a href="#">rules</a> ]	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[ <a href="#">error_arguments</a> ]	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 fields.

#### 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": "waf1.*",
        "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}
```

Deletes an event filter.

#### Related ONTAP commands

- `event filter delete`

#### Learn more

- [DOC /support/ems/filters/{name}](#)

#### Parameters

Name	Type	In	Required	Description
name	string	path	True	



## 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[ <a href="#">error_arguments</a> ]	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}

Retrieves an event filter.

### Related ONTAP commands

- `event filter show`

### Learn more

- [DOC /support/ems/filters/{name}](#)

### 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
<code>_links</code>	<code>_links</code>	
<code>name</code>	string	Filter name
<code>rules</code>	array[ <code>rules</code> ]	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	<a href="#">href</a>	

\_links

Name	Type	Description
related	<a href="#">href</a>	

message\_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
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	<a href="#">message_criteria</a>	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[ <a href="#">error_arguments</a> ]	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}

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

#### Learn more

- [DOC /support/ems/filters/{name}](#)

#### 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
_links	<a href="#">_links</a>	
name	string	Filter name
rules	array[ <a href="#">rules</a> ]	Array of event filter rules on which to match.

## Example request

```
{
  "_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"
  }
}
```

## Response

Status: 200, Ok

## Error

Status: Default

## ONTAP Error Response Codes

Error Code	Description
983088	The filter name provided is empty



Error Code	Description
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

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	<a href="#">href</a>	

\_links

Name	Type	Description
related	<a href="#">href</a>	

message\_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	
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	<a href="#">message_criteria</a>	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

#### ems\_filter

Name	Type	Description
_links	<a href="#">_links</a>	
name	string	Filter name
rules	array[ <a href="#">rules</a> ]	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[ <a href="#">error_arguments</a> ]	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 fields 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

Retrieves event filter rules.

### Related ONTAP commands

- `event filter show`

### Learn more

- [DOC /support/ems/filters/{name}/rules](#)

### 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.name_pattern	string	query	False	Filter by message_criteria.name_pattern
message_criteria.snmp_trap_types	string	query	False	Filter by message_criteria.snmp_trap_types
type	string	query	False	Filter by type

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.
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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

### Response

Status: 200, Ok

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">records</a> ]	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

\_links

Name	Type	Description
related	<a href="#">href</a>	

message\_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">_links</a>	

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	<a href="#">message_criteria</a>	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[ <a href="#">error_arguments</a> ]	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

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`

## Learn more

- [DOC /support/ems/filters/{name}/rules](#)

## Parameters

Name	Type	In	Required	Description
name	string	path	True	Filter name

## Request Body

Name	Type	Description
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">message_criteria</a>	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

## Example request

```
{
  "_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"
}
```

## Response

Status: 201, Created

Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">records</a> ]	

## 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

## 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

Error Code	Description
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

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	<a href="#">href</a>	

\_links

Name	Type	Description
related	<a href="#">href</a>	

message\_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
_links	<a href="#">_links</a>	
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
_links	<a href="#">_links</a>	
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	<a href="#">message_criteria</a>	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

\_links

Name	Type	Description
next	<a href="#">href</a>	
self	<a href="#">href</a>	

records

Rule for an event filter

Name	Type	Description
_links	<a href="#">_links</a>	
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	<a href="#">message_criteria</a>	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[ <a href="#">error_arguments</a> ]	Message arguments

Name	Type	Description
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 fields 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&sn
mp_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}
```

Deletes an event filter rule.

### Related ONTAP commands

- `event filter rule delete`

### Learn more

- [DOC /support/ems/filters/{name}/rules/{index}](#)

### Parameters

Name	Type	In	Required	Description
name	string	path	True	
index	string	path	True	

### 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

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[ <a href="#">error_arguments</a> ]	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}

Retrieves an event filter rule.

### Related ONTAP commands

- `event filter show`

### Learn more

- [DOC /support/ems/filters/{name}/rules/{index}](#)

### 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
<a href="#">_links</a>	<a href="#">_links</a>	
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	<a href="#">message_criteria</a>	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	<a href="#">href</a>	

\_links

Name	Type	Description
related	<a href="#">href</a>	

message\_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
_links	<a href="#">_links</a>	
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[ <a href="#">error_arguments</a> ]	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}

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`

#### Learn more

- [DOC /support/ems/filters/{name}/rules/{index}](#)

#### 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
<code>_links</code>	<a href="#">_links</a>	

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	<a href="#">message_criteria</a>	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

### Example request

```

{
  "_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"
}

```

### 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'
983146	An invalid value is provided for the property 'severities'
983147	The severity levels provided are not supported

Name	Type	Description
error	<a href="#">error</a>	

### 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	<a href="#">href</a>	

\_links

Name	Type	Description
related	<a href="#">href</a>	

message\_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
_links	<a href="#">_links</a>	
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
_links	<a href="#">_links</a>	
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	<a href="#">message_criteria</a>	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[ <a href="#">error_arguments</a> ]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

## Manage the 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

Retrieves the event catalog definitions.

#### Related ONTAP commands

- `event catalog show`

#### Learn more

- [DOC /support/ems/messages](#)

#### Parameters

Name	Type	In	Required	Description
description	string	query	False	Filter by description

Name	Type	In	Required	Description
snmp_trap_type	string	query	False	Filter by snmp_trap_type
deprecated	boolean	query	False	Filter by deprecated
corrective_action	string	query	False	Filter by corrective_action
severity	string	query	False	Filter by severity
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	The default is true for GET calls. When set to false, only the number of records is 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.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

#### Response

Status: 200, Ok



Name	Type	Description
_links	<a href="#">_links</a>	
num_records	integer	Number of records
records	array[ <a href="#">records</a> ]	

### 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	<a href="#">error</a>	

## 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	<a href="#">href</a>	
self	<a href="#">href</a>	

\_links

Name	Type	Description
self	<a href="#">href</a>	

records

Name	Type	Description
_links	<a href="#">_links</a>	
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

<b>Name</b>	<b>Type</b>	<b>Description</b>
arguments	array[ <a href="#">error_arguments</a> ]	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 © 2024 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.