Storage

ONTAP REST API reference

NetApp
February 26, 2025

This PDF was generated from https://docs.netapp.com/us-en/ontap-restapi-96/ontap/storage_overview.html on February 26, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Storage . . . 1
- Storage overview . . . 1
  - Overview . . . 1
- Retrieve or create a collection of storage aggregates . . . 1
  - Storage aggregates endpoint overview . . . 1
  - Retrieve a collection of aggregates for an entire cluster . . . 8
  - Create a collection of aggregates for an entire cluster . . . 27
- Retrieve a collection of cloud stores used by an aggregate . . . 42
  - Related ONTAP commands . . . 42
  - Parameters . . . 42
  - Response . . . 44
  - Error . . . 45
  - Definitions . . . 46
- Attach an object store to an aggregate or add a second object store as a mirror . . . 48
  - Required properties . . . 48
  - Recommended optional properties . . . 49
  - Default property values . . . 49
  - Related ONTAP commands . . . 49
  - Parameters . . . 49
  - Request Body . . . 50
  - Response . . . 51
  - Error . . . 52
  - Definitions . . . 52
- Remove a cloud target from an aggregate . . . 54
  - Related ONTAP commands . . . 54
  - Parameters . . . 55
  - Response . . . 55
  - Error . . . 55
  - Definitions . . . 56
- Retrieve the cloud store for an aggregate . . . 57
  - Related ONTAP commands . . . 58
  - Parameters . . . 58
  - Response . . . 58
  - Error . . . 59
  - Definitions . . . 60
- Update a cloud store . . . 62
  - Related ONTAP commands . . . 62
  - Parameters . . . 63
  - Request Body . . . 63
  - Response . . . 64
  - Error . . . 65
  - Definitions . . . 65
- Manage storage aggregate plexes . . . 67

Storage aggregates aggregate.uuid plexes endpoint overview . . . 67
Retrieve a collection of plexes for an aggregate . . . 70
Retrieve a plex specified by the aggregate UUID and plex name . . . 78
Delete an aggregate specified by the UUID . . . 85
Related ONTAP commands . . . 85
Learn more . . . 85
Parameters . . . 85
Response . . . 85
Error . . . 86
Definitions . . . 87
Retrieve an aggregate specified by the UUID . . . 88
Expensive properties . . . 89
Related ONTAP commands . . . 89
Learn more . . . 89
Parameters . . . 89
Response . . . 89
Error . . . 93
Definitions . . . 94
Update an aggregate specified by the UUID . . . 101
Related ONTAP commands . . . 101
Learn more . . . 101
Parameters . . . 101
Request Body . . . 102
Response . . . 105
Error . . . 105
Definitions . . . 108
Report cluster-wide storage details across different tiers . . . 116
Parameters . . . 116
Response . . . 116
Error . . . 117
Definitions . . . 118
Manage disks . . . 120
Storage disks endpoint overview . . . 120
Retrieve a collection of disks . . . 123
Retrieve a specific disk . . . 132
Create a clone of the file . . . 139
Request Body . . . 139
Response . . . 140
Error . . . 141
Definitions . . . 142
Manage FlexCache volumes . . . 144
Storage flexcache flexcaches endpoint overview . . . 144
Retrieve a FlexCache volume in the cluster . . . 149
Create a FlexCache volume in the cluster . . . 159
Delete a FlexCache volume . . . 168

Retrieve attributes of the FlexCache volume in the cluster . . . 171
Manage storage FlexCache origin volumes . . . 179
Storage flexcache origins endpoint overview . . . 179
Retrieve the origin of a FlexCache volume in the cluster . . . 184
Retrieve attributes of the origin of a FlexCache volume in the cluster . . . 192
Manage storage ports . . . 198
Storage ports endpoint overview . . . 198
Retrieve storage ports . . . 202
Retrieve a storage port . . . 209
Manage QoS policies . . . 214
Storage Qos policies endpoint overview . . . 214
Retrieve QoS policies . . . 216
Create a QoS policy . . . 224
Delete a QoS policy . . . 232
Retrieve a QoS policy . . . 235
Update a QoS policy . . . 241
Manage storage qtrees . . . 249
Storage qtrees endpoint overview . . . 249
Retrieve qtrees . . . 257
Create a qtree in a FlexVol or FlexGroup volume . . . 264
Delete a qtree . . . 271
Retrieve qtree properties . . . 274
Update properties for a qtree . . . 280
Storage quota . . . 287
Storage quota endpoint overview . . . 287
Manage storage quota reports . . . 288
Storage quota reports endpoint overview . . . 288
Retrieve the quota report records for all FlexVol and FlexGroup volumes . . . 305
Retrieve a specific quota report record . . . 316
Manage storage quota policy rules . . . 324
Storage quota rules endpoint overview . . . 325
Retrieve quota policy rules for all FlexVol and FlexGroup volumes . . . 341
Create a quota policy rule for a FlexVol or a FlexGroup volume . . . 351
Delete a quota policy rule . . . 364
Retrieve quota policy rule properties . . . 367
Update quota policy rule properties . . . 375
Manage shelves . . . 387
Storage shelves endpoint overview . . . 387
Retrieve shelves . . . 396
Retrieve a shelf . . . 407
Manage Snapshot copy policies . . . 414
Storage snapshot-policies endpoint overview . . . 414
Retrieve Snapshot copy policies . . . 420
Delete a Snapshot copy policy . . . 427
Retrieve Snapshot copy policy details . . . 429

Update a Snapshot copy policy . . . 434
Create a Snapshot copy policy . . . 440
Required properties . . . 440
Recommended optional properties . . . 440
Default property values . . . 440
Related ONTAP commands . . . 440
Learn more . . . 440
Request Body . . . 440
Response . . . 442
Error . . . 442
Definitions . . . 443
Manage storage volumes . . . 446
Storage volumes endpoint overview . . . 446
Retrieve volumes . . . 461
Create a volume on an SVM and storage aggregates . . . 510
Delete a volume . . . 552
Retrieve a volume . . . 554
Update volume attributes . . . 592
Retrieve historical performance metrics for a volume . . . 636
Parameters . . . 636
Response . . . 638
Error . . . 639
Definitions . . . 640
Manage volume Snapshot copies . . . 644
Storage volumes volume.uuid snapshots endpoint overview . . . 644
Retrieve volume Snapshot copies . . . 653
Create a volume Snapshot copy . . . 660
Delete a volume Snapshot copy . . . 667
Retrieve volume Snapshot copy details . . . 670
Update a volume Snapshot copy . . . 676

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 effeciency.

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-
b692-57d1cec874d4" -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
recommend	boolean	query	False	If set to 'true', it queries the system for the recommended optimal layout for creating new aggregates. The default setting is 'false'.
show_spares	boolean	query	False	If set to 'true', the spares object is returned instead of records to show the spare groups in the cluster. The default setting is 'false'.
data_encryption.software_encryption_enabled	boolean	query	False	Filter by data_encryption.software_encryption_enabled
data_encryption.drive_protection_enabled	boolean	query	False	Filter by data_encryption.drive_protection_enabled

Name	Type	In	Required	Description
state	string	query	False	Filter by state
space.efficiency_without_snapshots.ratio	number	query	False	Filter by space.efficiency_without_snapshots.ratio
space.efficiency_without_snapshots.logical_used	integer	query	False	Filter by space.efficiency_without_snapshots.logical_used
space.efficiency_without_snapshots.savings	integer	query	False	Filter by space.efficiency_without_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.efficiency.ratio	number	query	False	Filter by space.efficiency.ratio
space.efficiency.logical_used	integer	query	False	Filter by space.efficiency.logical_used
space.efficiency.savings	integer	query	False	Filter by space.efficiency.savings
space.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	_links	
error	error	
num_records	integer	Number of records
records	array[aggregate]	
spares	array[aggregate_spare]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "block_storage": {
        "hybrid_cache": {
          "disk_count": 6,
          "raid_type": "string",
          "size": 1612709888,
          "used": 26501122
        },
        "mirror": {
          "enabled": "",
          "state": "string"
        },
        "plexes": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
```

```
        },
        "name": "plex0"
      }
    ],
    "primary": {
      "checksum_style": "string",
      "disk_class": "performance",
      "disk_count": 8,
      "raid_size": 16,
      "raid_type": "string"
    }
  },
  "cloud_storage": {
    "stores": [
      {
        "cloud_store": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "store1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "used": 0
      }
    ]
  },
  "create_time": "2018-01-01 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": "string",
      "space": {
        "block_storage": {
          "available": 10156560384,
          "full_threshold_percent": 0,
          "inactive_user_data": 304448,
          "size": 10156769280,
          "used": 2088960
        },
        "cloud_storage": {
          "used": 402743264
        },
        "efficiency": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0
        },
        "efficiency_without_snapshots": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0
        },
        "footprint": 608896
      },
      "state": "string",
      "uuid": "string"
    }
  ],
  "spares": [
    {
      "checksum_style": "string",
      "disk_class": "solid_state",
      "layout_requirements": [
        {
          "aggregate_min_disks": 6,
          "raid_group": {
            "default": 16,
            "max": 28,
            "min": 5
          },
```

```
          "raid_type": "string"
        }
      ],
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "size": 10156769280,
      "syncmirror_pool": "string",
      "usable": 9
    }
  ]
}
```


Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
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	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

_links

Name	Type	Description
self	href	

hybrid_cache

Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.

Name	Type	Description
disk_count	integer	Number of disks used in the cache tier of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
enabled	boolean	Aggregate uses HDDs with SSDs as a cache
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
size	integer	Total usable space in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.
used	integer	Space used in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex_reference

Plex

Name	Type	Description
_links	_links	
name	string	

primary

Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

Name	Type	Description
checksum_style	string	The checksum style used by the aggregate.

Name	Type	Description
disk_class	string	The class of disks being used by the aggregate.
disk_count	integer	Number of disks used in the aggregate. This includes parity disks, but excludes disks in the hybrid cache.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type of the aggregate.

block_storage

Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.

Name	Type	Description
hybrid_cache	hybrid_cache	Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.
mirror	mirror	
plexes	array[plex_reference]	Plex reference for each plex in the aggregate.
primary	primary	Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

cloud_store

Cloud store

Name	Type	Description
_links	_links	
name	string	
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	_links	
name	string	
uuid	string	

node

Node where the aggregate currently resides.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

block_storage

Name	Type	Description
available	integer	Space available in bytes
full_threshold_percent	integer	The aggregate used percentage at which 'monitor.volume.full' EMS is generated.
inactive_user_data	integer	The size that is physically used in the block storage and has a cold temperature, in bytes. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the *fields* query parameter containing either block_storage.inactive_user_data or **.

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	block_storage	
cloud_storage	cloud_storage	
efficiency	efficiency	Storage efficiency
efficiency_without_snapshots	efficiency_without_snapshots	Storage efficiency that does not include the savings provided by Snapshot copies
footprint	integer	A summation of volume footprints (including volume guarantees), in bytes. This includes all of the volume footprints in the block_storage tier and the cloud_storage tier. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the *fields* query parameter containing either footprint or **.

aggregate

Name	Type	Description
_links	_links	
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation
data_encryption	data_encryption	

Name	Type	Description
dr_home_node	dr_home_node	Node where the aggregate belongs after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	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	string	Aggregate name
node	node	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	space	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

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	raid_group	
raid_type	string	RAID type

node

Node where the spares are assigned

Name	Type	Description
_links	_links	
name	string	
uuid	string	

aggregate_spare

Name	Type	Description
checksum_style	string	The checksum type that has been assigned to the spares
disk_class	string	Disk class of spares
layout_requirements	array[layout_requirement]	Available RAID protections and their restrictions
node	node	Node where the spares are assigned
size	integer	Usable size of each spare in bytes
syncmirror_pool	string	SyncMirror spare pool
usable	integer	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
disk_size	integer	query	False	If set, POST only selects disks of the specified size.

Request Body

Name	Type	Description
_links	_links	

Name	Type	Description
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate belongs after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	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	string	Aggregate name
node	node	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	space	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

Example request


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "raid_type": "string",
      "size": 1612709888,
      "used": 26501122
    },
    "mirror": {
      "enabled": "",
      "state": "string"
    },
    "plexes": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "plex0"
      }
    ],
    "primary": {
      "checksum_style": "string",
      "disk_class": "performance",
      "disk_count": 8,
      "raid_size": 16,
      "raid_type": "string"
    }
  },
  "cloud_storage": {
    "stores": [
      {
        "cloud_store": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
```

```
          "name": "store1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "used": 0
      }
    ]
  },
  "create_time": "2018-01-01 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": "string",
  "space": {
    "block_storage": {
      "available": 10156560384,
      "full_threshold_percent": 0,
      "inactive_user_data": 304448,
      "size": 10156769280,
      "used": 2088960
    },
    "cloud_storage": {
      "used": 402743264
    },
    "efficiency": {
      "logical_used": 0,
      "ratio": 0,
```

```
      "savings": 0
    },
    "efficiency_without_snapshots": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "footprint": 608896
  },
  "state": "string",
  "uuid": "string"
}
```


Response

```
Status: 202, Accepted
```

Name	Type	Description
job	job_link	

Example response


```
{
  "job": {
    "_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.
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	href	

hybrid_cache

Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.

Name	Type	Description
disk_count	integer	Number of disks used in the cache tier of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
enabled	boolean	Aggregate uses HDDs with SSDs as a cache
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
size	integer	Total usable space in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.
used	integer	Space used in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex_reference

Plex

Name	Type	Description
_links	_links	
name	string	

primary

Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

Name	Type	Description
checksum_style	string	The checksum style used by the aggregate.
disk_class	string	The class of disks being used by the aggregate.
disk_count	integer	Number of disks used in the aggregate. This includes parity disks, but excludes disks in the hybrid cache.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type of the aggregate.

block_storage

Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.

Name	Type	Description
hybrid_cache	hybrid_cache	Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.
mirror	mirror	
plexes	array[plex_reference]	Plex reference for each plex in the aggregate.

Name	Type	Description
primary	primary	Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

cloud_store

Cloud store

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cloud_storage_tier

Name	Type	Description
cloud_store	cloud_store	Cloud store
used	integer	Capacity used in bytes in the cloud store by this aggregate. This is a cached value calculated every 5 minutes.

cloud_storage

Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.

Name	Type	Description
attach_eligible	boolean	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	_links	
name	string	
uuid	string	

node

Node where the aggregate currently resides.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

block_storage

Name	Type	Description
available	integer	Space available in bytes

Name	Type	Description
full_threshold_percent	integer	The aggregate used percentage at which 'monitor.volume.full' EMS is generated.
inactive_user_data	integer	The size that is physically used in the block storage and has a cold temperature, in bytes. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the *fields* query parameter containing either block_storage.inactive_user_data or **.
size	integer	Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
used	integer	Space used or reserved in bytes. Includes volume guarantees and aggregate metadata.

cloud_storage

Name	Type	Description
used	integer	Used space in bytes in the cloud store. Only applicable for 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	block_storage	
cloud_storage	cloud_storage	
efficiency	efficiency	Storage efficiency
efficiency_without_snapshots	efficiency_without_snapshots	Storage efficiency that does not include the savings provided by Snapshot copies
footprint	integer	A summation of volume footprints (including volume guarantees), in bytes. This includes all of the volume footprints in the block_storage tier and the cloud_storage tier. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the *fields* query parameter containing either footprint or **.

aggregate

Name	Type	Description
_links	_links	
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate belongs after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	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	string	Aggregate name
node	node	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	space	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a 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

Name	Type	In	Required	Description
unreclaimed_space_threshold	integer	query	False	Filter by unreclaimed_space_threshold
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	_links	
num_records	integer	Number of records
records	array[cloud_store]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "availability": "string",
      "target": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "target1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "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	href	
self	href	

_links

Name	Type	Description
self	href	

target

Cloud target

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cloud_store

Name	Type	Description
_links	_links	
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a cloud mirrored composite aggregate.
target	target	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[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Attach an object store to an aggregate or add a second object store as a mirror

POST `/storage/aggregates/{aggregate.uuid}/cloud-stores`

Attaches an object store to an aggregate, or adds a second object store as a mirror.

Required properties

- `target.uuid` or `target.name` - UUID or name of the cloud target.

Recommended optional properties

- `primary` - *true* if the object store is primary or *false* if it is a mirror.
- `allow_flexgroups` - Allow attaching object store to an aggregate containing FlexGroup constituents.
- `check_only` - Validate only and do not add the cloud store.

Default property values

- `primary` - *true*
- `allow_flexgroups` - *false*
- `check_only` - *false*

Related ONTAP commands

- `storage aggregate object-store attach`
- `storage aggregate object-store mirror`

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
allow_flexgroups	boolean	query	False	This optional parameter allows attaching object store to an aggregate containing FlexGroup constituents. The default value is false. Mixing FabricPools and non-FabricPools within a FlexGroup is not recommended. All aggregates hosting constituents of a FlexGroup should be attached to the object store.
check_only	boolean	query	False	Validate only and do not add the cloud store.

Request Body

Name	Type	Description
_links	_links	
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a cloud mirrored composite aggregate.
target	target	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": "string",
  "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	job_link	

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	href	

target

Cloud target

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cloud_store

Name	Type	Description
_links	_links	
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a cloud mirrored composite aggregate.
target	target	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	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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	job_link	

Example response


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


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve the cloud store for an aggregate

GET `/storage/aggregates/{aggregate.uuid}/cloud-stores/{target.uuid}`

Retrieves the cloud store for the aggregate using the specified cloud target UUID.

Related ONTAP commands

- `storage aggregate object-store show`

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
target.uuid	string	path	True	Cloud target UUID
fields	array[string]	query	False	Specify the fields to return.

Response

```
Status: 200, Ok
```

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

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "availability": "string",
      "target": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "target1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "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	href	
self	href	

_links

Name	Type	Description
self	href	

target

Cloud target

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cloud_store

Name	Type	Description
_links	_links	
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a cloud mirrored composite aggregate.
target	target	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[error_arguments]	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
_links	_links	
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a cloud mirrored composite aggregate.
target	target	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": "string",
  "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	job_link	

Example response


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


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

target

Cloud target

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cloud_store

Name	Type	Description
_links	_links	
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a cloud mirrored composite aggregate.
target	target	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	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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
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.

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	_links	
error	error	
num_records	integer	Number of records
records	array[plex]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "records": [
    {
      "aggregate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "plex0",
      "pool": "string",
      "raid_groups": [
        {
          "disks": [
            {
              "disk": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                },
                "name": "1.0.1"
```

```
              },
              "position": "string",
              "state": "string",
              "type": "ssd",
              "usable_size": 947912704
            }
          ],
          "name": "rg0",
          "recomputing_parity": {
            "percent": 10
          },
          "reconstruct": {
            "percent": 10
          }
        }
      ],
      "resync": {
        "level": "string",
        "percent": 10
      },
      "state": "string"
    }
  ]
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

_links

Name	Type	Description
self	href	

aggregate

Aggregate

Name	Type	Description
_links	_links	
name	string	

Name	Type	Description
uuid	string	

disk

Disk

Name	Type	Description
_links	_links	
name	string	

raid_group_disk

Name	Type	Description
disk	disk	Disk
position	string	The position of the disk within the RAID group.
state	string	The state of the disk within the RAID group.
type	string	Disk interface type
usable_size	integer	Size in bytes that is usable by the aggregate.

recomputing_parity

Name	Type	Description
active	boolean	RAID group is recomputing parity
percent	integer	Recomputing parity percentage

reconstruct

Name	Type	Description
active	boolean	One or more disks in this RAID group are being reconstructed.
percent	integer	Reconstruct percentage

raid_group

Name	Type	Description
cache_tier	boolean	RAID group is a cache tier
degraded	boolean	RAID group is degraded. A RAID group is degraded when at least one disk from that group has failed or is offline.
disks	array[raid_group_disk]	
name	string	RAID group name
recomputing_parity	recomputing_parity	
reconstruct	reconstruct	

resync

Name	Type	Description
active	boolean	Plex is being resynchronized to its mirrored plex
level	string	Plex resyncing level
percent	integer	Plex resyncing percentage

plex

Name	Type	Description
aggregate	aggregate	Aggregate
name	string	Plex name
online	boolean	Plex is online
pool	string	SyncMirror pool assignment
raid_groups	array[raid_group]	
resync	resync	
state	string	Plex state

Retrieve a plex specified by the aggregate UUID and plex name

GET `/storage/aggregates/{aggregate.uuid}/plexes/{name}`

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	aggregate	Aggregate
name	string	Plex name
online	boolean	Plex is online
pool	string	SyncMirror pool assignment
raid_groups	array[raid_group]	
resync	resync	
state	string	Plex state

Example response

```
{
  "aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "plex0",
  "pool": "string",
  "raid_groups": [
    {
      "disks": [
        {
          "disk": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "1.0.1"
          },
          "position": "string",
          "state": "string",
          "type": "ssd",
          "usable_size": 947912704
        }
      ],
      "name": "rg0",
      "recomputing_parity": {
        "percent": 10
      },
      "reconstruct": {
        "percent": 10
      }
    }
  ],
  "resync": {
    "level": "string",
    "percent": 10
  },
  "state": "string"
}
```

Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregate

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

disk

Disk

Name	Type	Description
_links	_links	
name	string	

raid_group_disk

Name	Type	Description
disk	disk	Disk
position	string	The position of the disk within the RAID group.
state	string	The state of the disk within the RAID group.
type	string	Disk interface type
usable_size	integer	Size in bytes that is usable by the aggregate.

recomputing_parity

Name	Type	Description
active	boolean	RAID group is recomputing parity
percent	integer	Recomputing parity percentage

reconstruct

Name	Type	Description
active	boolean	One or more disks in this RAID group are being reconstructed.
percent	integer	Reconstruct percentage

raid_group

Name	Type	Description
cache_tier	boolean	RAID group is a cache tier
degraded	boolean	RAID group is degraded. A RAID group is degraded when at least one disk from that group has failed or is offline.
disks	array[raid_group_disk]	
name	string	RAID group name
recomputing_parity	recomputing_parity	
reconstruct	reconstruct	

resync

Name	Type	Description
active	boolean	Plex is being resynchronized to its mirrored plex
level	string	Plex resyncing level
percent	integer	Plex resyncing percentage

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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	job_link	

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	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve an aggregate specified by the UUID

GET `/storage/aggregates/{uuid}`

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
_links	_links	
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation

Name	Type	Description
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate belongs after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	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	string	Aggregate name
node	node	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	space	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "raid_type": "string",
      "size": 1612709888,
      "used": 26501122
    },
    "mirror": {
      "enabled": "",
      "state": "string"
    },
    "plexes": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "plex0"
      }
    ],
    "primary": {
      "checksum_style": "string",
      "disk_class": "performance",
      "disk_count": 8,
      "raid_size": 16,
      "raid_type": "string"
    }
  },
  "cloud_storage": {
    "stores": [
      {
        "cloud_store": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
```

```
          "name": "store1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "used": 0
      }
    ]
  },
  "create_time": "2018-01-01 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": "string",
  "space": {
    "block_storage": {
      "available": 10156560384,
      "full_threshold_percent": 0,
      "inactive_user_data": 304448,
      "size": 10156769280,
      "used": 2088960
    },
    "cloud_storage": {
      "used": 402743264
    },
    "efficiency": {
      "logical_used": 0,
      "ratio": 0,
```

```
      "savings": 0
    },
    "efficiency_without_snapshots": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "footprint": 608896
  },
  "state": "string",
  "uuid": "string"
}
```


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	href	

hybrid_cache

Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.

Name	Type	Description
disk_count	integer	Number of disks used in the cache tier of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
enabled	boolean	Aggregate uses HDDs with SSDs as a cache
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
size	integer	Total usable space in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.
used	integer	Space used in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex_reference

Plex

Name	Type	Description
_links	_links	
name	string	

primary

Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

Name	Type	Description
checksum_style	string	The checksum style used by the aggregate.
disk_class	string	The class of disks being used by the aggregate.
disk_count	integer	Number of disks used in the aggregate. This includes parity disks, but excludes disks in the hybrid cache.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type of the aggregate.

block_storage

Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.

Name	Type	Description
hybrid_cache	hybrid_cache	Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.
mirror	mirror	
plexes	array[plex_reference]	Plex reference for each plex in the aggregate.

Name	Type	Description
primary	primary	Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

cloud_store

Cloud store

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cloud_storage_tier

Name	Type	Description
cloud_store	cloud_store	Cloud store
used	integer	Capacity used in bytes in the cloud store by this aggregate. This is a cached value calculated every 5 minutes.

cloud_storage

Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.

Name	Type	Description
attach_eligible	boolean	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	_links	
name	string	
uuid	string	

node

Node where the aggregate currently resides.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

block_storage

Name	Type	Description
available	integer	Space available in bytes

Name	Type	Description
full_threshold_percent	integer	The aggregate used percentage at which 'monitor.volume.full' EMS is generated.
inactive_user_data	integer	The size that is physically used in the block storage and has a cold temperature, in bytes. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the *fields* query parameter containing either block_storage.inactive_user_data or **.
size	integer	Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
used	integer	Space used or reserved in bytes. Includes volume guarantees and aggregate metadata.

cloud_storage

Name	Type	Description
used	integer	Used space in bytes in the cloud store. Only applicable for 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	block_storage	
cloud_storage	cloud_storage	
efficiency	efficiency	Storage efficiency
efficiency_without_snapshots	efficiency_without_snapshots	Storage efficiency that does not include the savings provided by Snapshot copies
footprint	integer	A summation of volume footprints (including volume guarantees), in bytes. This includes all of the volume footprints in the block_storage tier and the cloud_storage tier. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the *fields* 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[error_arguments]	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	_links	
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate belongs after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	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	string	Aggregate name
node	node	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	space	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

Example request


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "raid_type": "string",
      "size": 1612709888,
      "used": 26501122
    },
    "mirror": {
      "enabled": "",
      "state": "string"
    },
    "plexes": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "plex0"
      }
    ],
    "primary": {
      "checksum_style": "string",
      "disk_class": "performance",
      "disk_count": 8,
      "raid_size": 16,
      "raid_type": "string"
    }
  },
  "cloud_storage": {
    "stores": [
      {
        "cloud_store": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
```

```
          "name": "store1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "used": 0
      }
    ]
  },
  "create_time": "2018-01-01 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": "string",
  "space": {
    "block_storage": {
      "available": 10156560384,
      "full_threshold_percent": 0,
      "inactive_user_data": 304448,
      "size": 10156769280,
      "used": 2088960
    },
    "cloud_storage": {
      "used": 402743264
    },
    "efficiency": {
      "logical_used": 0,
      "ratio": 0,
```

```
      "savings": 0
    },
    "efficiency_without_snapshots": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "footprint": 608896
  },
  "state": "string",
  "uuid": "string"
}
```


Response

```
Status: 202, Accepted
```

Name	Type	Description
job	job_link	

Example response


```
{
  "job": {
    "_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.
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.

Error Code	Description
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.
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	href	

hybrid_cache

Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.

Name	Type	Description
disk_count	integer	Number of disks used in the cache tier of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
enabled	boolean	Aggregate uses HDDs with SSDs as a cache
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
size	integer	Total usable space in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.
used	integer	Space used in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex_reference

Plex

Name	Type	Description
_links	_links	
name	string	

primary

Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

Name	Type	Description
checksum_style	string	The checksum style used by the aggregate.
disk_class	string	The class of disks being used by the aggregate.
disk_count	integer	Number of disks used in the aggregate. This includes parity disks, but excludes disks in the hybrid cache.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type of the aggregate.

block_storage

Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.

Name	Type	Description
hybrid_cache	hybrid_cache	Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.
mirror	mirror	
plexes	array[plex_reference]	Plex reference for each plex in the aggregate.

Name	Type	Description
primary	primary	Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

cloud_store

Cloud store

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cloud_storage_tier

Name	Type	Description
cloud_store	cloud_store	Cloud store
used	integer	Capacity used in bytes in the cloud store by this aggregate. This is a cached value calculated every 5 minutes.

cloud_storage

Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.

Name	Type	Description
attach_eligible	boolean	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	_links	
name	string	
uuid	string	

node

Node where the aggregate currently resides.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

block_storage

Name	Type	Description
available	integer	Space available in bytes

Name	Type	Description
full_threshold_percent	integer	The aggregate used percentage at which 'monitor.volume.full' EMS is generated.
inactive_user_data	integer	The size that is physically used in the block storage and has a cold temperature, in bytes. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the *fields* query parameter containing either block_storage.inactive_user_data or **.
size	integer	Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
used	integer	Space used or reserved in bytes. Includes volume guarantees and aggregate metadata.

cloud_storage

Name	Type	Description
used	integer	Used space in bytes in the cloud store. Only applicable for 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	block_storage	
cloud_storage	cloud_storage	
efficiency	efficiency	Storage efficiency
efficiency_without_snapshots	efficiency_without_snapshots	Storage efficiency that does not include the savings provided by Snapshot copies
footprint	integer	A summation of volume footprints (including volume guarantees), in bytes. This includes all of the volume footprints in the block_storage tier and the cloud_storage tier. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the *fields* query parameter containing either footprint or **.

aggregate

Name	Type	Description
_links	_links	
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate belongs after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	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	string	Aggregate name
node	node	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	space	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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": "string"
      }
    ]
  },
  "cloud_storage": {
    "used": 0
  },
  "efficiency": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  }
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

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

Manage disks

Storage disks endpoint overview

Retrieving storage disk information

The storage disk GET API retrieves all of the disks in the cluster.

Examples

1) Retrieve a list of disks from the cluster

The following example shows the response with a list of disks in the cluster:

```
# The API:
/api/storage/disks

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

# The response:
{
"records": [
  {
    "name": "1.24.4",
```

```
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.4"
      }
    }
  },
  {
    "name": "1.24.3",
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.3"
      }
    }
  },
  {
    "name": "1.24.5",
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.5"
      }
    }
  },
  {
    "name": "1.24.0",
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.0"
      }
    }
  },
  {
    "name": "1.24.2",
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.2"
      }
    }
  },
  {
    "name": "1.24.1",
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.1"
      }
    }
  }
```

```
],
"num_records": 6,
"_links": {
  "self": {
    "href": "/api/storage/disks"
  }
}
}
```

2) Retrieve a specific disk from the cluster

The following example shows the response of the requested disk. If there is no disk with the requested name, an error is returned.

```
# The API:
/api/storage/disks/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/disks/1.24.3" -H "accept:
application/hal+json"

# The response:
{
"name": "1.24.3",
"uid":
"50000394:0808AA88:00000000:00000000:00000000:00000000:00000000:00000000:0
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.uid	string	query	False	Filter by shelf.uid
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	_links	
num_records	integer	Number of records
records	array[disk]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "aggregates": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "bay": 1,
      "class": "solid_state",
      "container_type": "spare",
      "dr_node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "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	href	
self	href	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

dr_node

Name	Type	Description
name	string	
uuid	string	

drawer

Name	Type	Description
id	integer	
slot	integer	

home_node

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	
uuid	string	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

shelf_reference

Shelf

Name	Type	Description
_links	_links	
uid	string	

disk

Name	Type	Description
aggregates	array[aggregates]	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	dr_node	
drawer	drawer	
firmware_version	string	
home_node	home_node	
model	string	
name	string	Cluster-wide disk name
node	node	
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	shelf_reference	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[error_arguments]	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[aggregates]	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	dr_node	
drawer	drawer	
firmware_version	string	
home_node	home_node	
model	string	
name	string	Cluster-wide disk name
node	node	
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	shelf_reference	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:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000",
  "usable_size": 959934889984,
  "vendor": "NETAPP"
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

dr_node

Name	Type	Description
name	string	
uuid	string	

drawer

Name	Type	Description
id	integer	
slot	integer	

home_node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

shelf_reference

Shelf

Name	Type	Description
_links	_links	
uid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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.
destination_path	string	Relative path of the clone/destination file in the volume.
is_backup	boolean	Mark clone file for backup.
overwrite_destination	boolean	Destination file gets overwritten.
range	array[string]	List of block ranges for sub-file cloning in the format "source-file-block-number:destination-file-block-number:block-count"
source_path	string	Relative path of the source file in the volume.
volume	volume	

Example request

```
{
  "destination_path": "dest_file1, dir1/dest_file2",
  "range": [
    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	href	

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

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

Manage FlexCache volumes

Storage flexcache flexcaches endpoint overview

Overview

FlexCache is a persistent cache of an origin volume. An origin volume can only be a FlexVol while a FlexCache is always a FlexGroup.

The following relationship configurations are supported:

– Intra-Vserver where FlexCache and the corresponding origin volume reside in the same Vserver.

– Cross-Vserver but intra-cluster where FlexCache and the origin volume reside in the same cluster but belong to different Vservers.

– Cross-cluster where FlexCache and the origin volume reside in different clusters.

FlexCache supports fan-out and more than one FlexCache can be created from one origin volume. This API retrieves and manages FlexCache configurations in the cache cluster.

FlexCache APIs

The following APIs can be used to perform operations related with FlexCache:

– GET /api/storage/flexcache/flexcaches

– GET /api/storage/flexcache/flexcaches/{uuid}

– POST /api/storage/flexcache/flexcaches

– DELETE /api/storage/flexcache/flexcaches/{uuid}

Examples

Creating a FlexCache

The POST request is used to create a FlexCache.

```
# The API:
/api/storage/flexcache/flexcaches

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/flexcache/flexcaches" -H
"accept: application/json" -H  "Content-Type: application/json" -d "{
\"aggregates\": [ { \"name\": \"aggr_1\" } ], \"name\": \"fc_333\",
\"origins\": [ {  \"svm\": { \"name\": \"vs_3\"  }, \"volume\": {
\"name\": \"vol_o1\" } } ], \"svm\": { \"name\": \"vs_1\" } }"

# The response:
{
"job": {
  "uuid": "e751dd5d-0f3c-11e9-8b2b-0050568e0b79",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e751dd5d-0f3c-11e9-8b2b-0050568e0b79"
    }
  }
}
}
```

Retrieving FlexCache attributes

The GET request is used to retrieve FlexCache attributes. The object includes a large set of fields which can be expensive to retrieve. Most notably, the fields size, 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_o1",
        "uuid": "2bc957dd-2617-4afb-8d2f-66ac6070d313"
      },
      "svm": {
        "name": "vs_3",
        "uuid": "8aa2cd28-0e92-11e9-b391-0050568e4115"
      },
      "cluster": {
        "name": "node2",
        "uuid": "50733f81-0e90-11e9-b391-0050568e4115"
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-
0050568e0b79"
    }
  }
}
```

Deleting a FlexCache

The DELETE request is used to delete a FlexCache.

```
# The API:
/api/storage/flexcache/flexcaches

# The call:
curl -X DELETE "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-0050568e0b79"
-H  "accept: application/json"

# The response:
{
  "job": {
    "uuid": "e17994f2-0f3e-11e9-8b2b-0050568e0b79",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/e17994f2-0f3e-11e9-8b2b-0050568e0b79"
      }
    }
  }
}
```

Retrieve a FlexCache volume in the cluster

GET `/storage/flexcache/flexcaches`

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 aggregrate 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

Name	Type	In	Required	Description
origins.create_time	string	query	False	Filter by origins.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.
max_records	integer	query	False	Limit the number of records returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	return_records	boolean	query	False

Response

```
Status: 200, Ok
```

Name	Type	Description
_links	_links	
num_records	integer	Number of records

Name	Type	Description
records	array[flexcache]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "aggregates": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "name": "vol1",
      "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": "string",
          "svm": {
```

```
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "svm1",
            "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
          },
          "volume": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "volume1",
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
          }
        }
      ],
      "path": "/user/my_fc",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster managerment IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	SVM, applies only to SVM-scoped objects.
volume	volume	

svm

FlexCache SVM

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

flexcache

Defines the cache endpoint of FlexCache.

Name	Type	Description
_links	_links	
aggregates	array[aggregates]	
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
name	string	FlexCache name
origins	array[flexcache_relationship]	
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

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	_links	
aggregates	array[aggregates]	

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[flexcache_relationship]	
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM
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": "vol1",
  "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": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
```

```
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "path": "/user/my_fc",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


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
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	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	

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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster managerment IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	SVM, applies only to SVM-scoped objects.
volume	volume	

svm

FlexCache SVM

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

flexcache

Defines the cache endpoint of FlexCache.

Name	Type	Description
_links	_links	
aggregates	array[aggregates]	
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
name	string	FlexCache name
origins	array[flexcache_relationship]	
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.

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

Name	Type	Description
message	string	Message argument

error

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

Delete a 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	job_link	

Example response


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


Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
66846879	The specified volume UUID is not a FlexCache volume
66846731	Failed to delete the FlexCache volume
524546	Failed to delete the FlexCache volume because the FlexCache volume is not unmounted

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve attributes of the FlexCache volume in the cluster

GET `/storage/flexcache/flexcaches/{uuid}`

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 aggregrate of FlexCache volume.
- `path` - Fully-qualified path of the owning SVM's namespace where the FlexCache is mounted.

Related ONTAP commands

- `volume flexcache show`

Learn more

- DOC /storage/flexcache/flexcaches

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Unique identifier of FlexCache.
fields	array[string]	query	False	Specify the fields to return.

Response

```
Status: 200, Ok
```

Name	Type	Description
_links	_links	
aggregates	array[aggregates]	
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.

Name	Type	Description
name	string	FlexCache name
origins	array[flexcache_relationship]	
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM
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": "vol1",
  "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": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
```

```
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "path": "/user/my_fc",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	

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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster managerment IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	SVM, applies only to SVM-scoped objects.
volume	volume	

svm

FlexCache SVM

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

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

Manage storage FlexCache origin volumes

Storage flexcache origins endpoint overview

Overview

FlexCache is a persistent cache of an origin volume. An origin volume can only be a FlexVol while a FlexCache is always a FlexGroup.

The following relationship configurations are supported:

– Intra-Vserver where FlexCache and the corresponding origin volume reside in the same Vserver.

– Cross-Vserver but intra-cluster where FlexCache and the origin volume reside in the same cluster but belong to different Vservers.

– Cross-cluster where FlexCache and the origin volume reside in different clusters.

FlexCache supports fan-out and more than one FlexCache can be created from one origin volume. This API retrieves the origin of FlexCache onfigurations in the origin cluster.

FlexCache APIs

The following APIs can be used to perform operations related to the origin of a FlexCache:

– GET /api/storage/flexcache/origins

– GET /api/storage/flexcache/origins/{uuid}

Examples

Retrieving origins of FlexCache attributes

The GET request is used to retrieve the origins of FlexCache attributes.

```
# The API:
/api/storage/flexcache/origins

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/flexcache/origins?" -H
"accept: application/json"

# The response:
{
  "records": [
    {
      "uuid": "2bc957dd-2617-4afb-8d2f-66ac6070d313",
      "name": "vol_o1",
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/origins/2bc957dd-2617-4afb-8d2f-
66ac6070d313"
        }
      }
    },
    {
      "uuid": "80fcaee4-0dc2-488b-afb8-86d28a34cda8",
      "name": "vol_1",
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/origins/80fcaee4-0dc2-488b-afb8-
86d28a34cda8"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/storage/flexcache/origins?"
    }
  }
}
```

Retrieving the attributes of an origin volume

The GET request is used to retrieve the attributes of an origin volume.

```
# The API:
/api/storage/flexcache/origins/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/flexcache/origins/80fcaee4-
0dc2-488b-afb8-86d28a34cda8" -H  "accept: application/json"

# The response:
{
  "uuid": "80fcaee4-0dc2-488b-afb8-86d28a34cda8",
  "name": "vol_1",
  "svm": {
    "name": "vs_3",
    "uuid": "8aa2cd28-0e92-11e9-b391-0050568e4115"
  },
  "flexcaches": [
    {
      "ip_address": "10.140.103.183",
      "create_time": "2019-01-02T19:27:22+05:30",
      "volume": {
        "name": "fc_42",
        "uuid": "4e7f9d49-0e96-11e9-aed0-0050568eddbe"
      },
      "svm": {
        "name": "vs_1_4",
        "uuid": "36f68322-0e93-11e9-aed0-0050568eddbe"
      },
      "cluster": {
        "name": "node4",
        "uuid": "c32f16b8-0e90-11e9-aed0-0050568eddbe"
      }
    },
    {
      "ip_address": "10.140.103.183",
      "create_time": "2019-01-02T21:08:34+05:30",
      "volume": {
        "name": "fc_421",
        "uuid": "71ee8f36-0ea4-11e9-aed0-0050568eddbe"
      },
      "svm": {
        "name": "vs_1_4",
        "uuid": "36f68322-0e93-11e9-aed0-0050568eddbe"
      },
      "cluster": {
        "name": "node4",
        "uuid": "c32f16b8-0e90-11e9-aed0-0050568eddbe"
```

```
      }
    },
    {
      "ip_address": "10.140.103.183",
      "create_time": "2019-01-03T11:14:38+05:30",
      "volume": {
        "name": "fc_422"
      },
      "svm": {
        "name": "vs_1_4",
        "uuid": "36f68322-0e93-11e9-aed0-0050568eddbe"
      },
      "cluster": {
        "name": "node4",
        "uuid": "c32f16b8-0e90-11e9-aed0-0050568eddbe"
      }
    },
    {
      "ip_address": "10.140.103.179",
      "size": 4294967296,
      "create_time": "2019-01-02T19:24:14+05:30",
      "state": "online",
      "volume": {
        "name": "fc_32",
        "uuid": "ddb42bbc-0e95-11e9-8180-0050568e0b79"
      },
      "svm": {
        "name": "vs_1",
        "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
      },
      "cluster": {
        "name": "node3",
        "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
      }
    },
    {
      "ip_address": "10.140.103.179",
      "size": 4294967296,
      "create_time": "2019-01-02T21:07:23+05:30",
      "state": "online",
      "volume": {
        "name": "fc_321",
        "uuid": "47902654-0ea4-11e9-8180-0050568e0b79"
      },
      "svm": {
        "name": "vs_1",
```

```
        "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
      },
      "cluster": {
        "name": "node3",
        "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
      }
    },
    {
      "ip_address": "10.140.103.179",
      "size": 4294967296,
      "create_time": "2019-01-03T00:11:38+05:30",
      "state": "online",
      "volume": {
        "name": "fc_322",
        "uuid": "04d5e07b-0ebe-11e9-8180-0050568e0b79"
      },
      "svm": {
        "name": "vs_1",
        "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
      },
      "cluster": {
        "name": "node3",
        "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
      }
    },
    {
      "ip_address": "10.140.103.179",
      "size": 4294967296,
      "create_time": "2019-01-03T00:14:52+05:30",
      "state": "online",
      "volume": {
        "name": "fc_323",
        "uuid": "77e911ff-0ebe-11e9-8180-0050568e0b79"
      },
      "svm": {
        "name": "vs_1",
        "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
      },
      "cluster": {
        "name": "node3",
        "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
      }
    }
  ],
  "_links": {
    "self": {
```

```
      "href": "/api/storage/flexcache/origins/80fcaee4-0dc2-488b-afb8-
86d28a34cda8"
    }
  }
}
```

Retrieve the origin of a FlexCache volume in the cluster

GET `/storage/flexcache/origins`

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_address	string	query	False	Filter by flexcaches.ip_address

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	_links	
num_records	integer	Number of records
records	array[flexcache_origin]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "flexcaches": [
        {
          "cluster": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "cluster1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "create_time": "2018-06-04 19:00:00 UTC",
          "ip_address": "10.10.10.7",
          "size": 0,
          "state": "string",
          "svm": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "svm1",
            "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
          },
          "volume": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
```

```
              }
            },
            "name": "volume1",
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
          }
        }
      ],
      "name": "vol1, vol_2",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563512"
    }
  ]
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster managerment IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	SVM, applies only to SVM-scoped objects.
volume	volume	

svm

Origin volume SVM

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

flexcache_origin

Defines the origin endpoint of FlexCache.

Name	Type	Description
_links	_links	
flexcaches	array[flexcache_relationship]	

Name	Type	Description
name	string	Origin volume name
svm	svm	Origin volume SVM
uuid	string	Origin volume UUID. Unique identifier for origin of FlexCache.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve attributes of the origin of a FlexCache volume in the cluster

GET `/storage/flexcache/origins/{uuid}`

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
_links	_links	
flexcaches	array[flexcache_relationship]	
name	string	Origin volume name
svm	svm	Origin volume SVM
uuid	string	Origin volume UUID. Unique identifier for origin of FlexCache.

Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "flexcaches": [
    {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "create_time": "2018-06-04 19:00:00 UTC",
      "ip_address": "10.10.10.7",
      "size": 0,
      "state": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "name": "vol1, vol_2",
  "svm": {
    "_links": {
```

```
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563512"
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster managerment IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	SVM, applies only to SVM-scoped objects.
volume	volume	

svm

Origin volume SVM

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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	_links	

Name	Type	Description
num_records	integer	Number of records
records	array[storage_port]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "board_name": "string",
      "cable": {
        "identifier": "500a0980000b6c3f-50000d1703544b80",
        "length": "2m",
        "part_number": "112-00431+A0",
        "serial_number": 616930439
      },
      "description": "SAS Host Adapter 2a (PMC-Sierra PM8072 rev. C)",
      "error": {
        "corrective_action": "string",
        "message": "string"
      },
      "mac_address": "string",
      "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	href	
self	href	

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

error

Name	Type	Description
corrective_action	string	Error corrective action
message	string	Error message

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_port

Name	Type	Description
board_name	string	
cable	cable	
description	string	
error	error	
mac_address	string	
name	string	
node	node	
part_number	string	
serial_number	string	
speed	number	Operational port speed in Gbps
state	string	
wwn	string	World Wide Name

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a storage port

GET `/storage/ports/{node.uuid}/{name}`

Retrieves a specific storage port.

Related ONTAP commands

- `storage port show`

Learn more

- DOC /storage/ports

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Node UUID
name	string	path	True	Port name
fields	array[string]	query	False	Specify the fields to return.

Response

```
Status: 200, Ok
```

Name	Type	Description
board_name	string	
cable	cable	
description	string	
error	error	
mac_address	string	
name	string	
node	node	
part_number	string	
serial_number	string	
speed	number	Operational port speed in Gbps
state	string	
wwn	string	World Wide Name

Example response


```
{
  "board_name": "string",
  "cable": {
    "identifier": "500a0980000b6c3f-50000d1703544b80",
    "length": "2m",
    "part_number": "112-00431+A0",
    "serial_number": 616930439
  },
  "description": "SAS Host Adapter 2a (PMC-Sierra PM8072 rev. C)",
  "error": {
    "corrective_action": "string",
    "message": "string"
  },
  "mac_address": "string",
  "name": "2a",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "part_number": "111-03801",
  "serial_number": "7A2463CC45B",
  "speed": 6,
  "state": "online",
  "wwn": "50000d1703544b80"
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

error

Name	Type	Description
corrective_action	string	Error corrective action
message	string	Error message

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage QoS policies

Storage Qos policies endpoint overview

Quality of Service Configuration

A QoS policy defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy is associated. There are two types of policies that can be configured, fixed which 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 exisiting 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	_links	
error	error	
num_records	integer	Number of records
records	array[qos_policy]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "extreme",
      "object_count": 0,
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

_links

Name	Type	Description
self	href	

adaptive

Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

Name	Type	Description
absolute_min_iops	integer	Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
expected_iops	integer	Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
peak_iops	integer	Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size.

fixed

QoS policy-groups define a fixed service level objective (SLO) for a storage object.

Name	Type	Description
capacity_shared	boolean	Specifies whether the capacities are shared across all objects that use this QoS policy-group. Default is false.
max_throughput_iops	integer	Maximum throughput defined by this policy. It is specified in terms of IOPS. 0 means no maximum throughput is enforced.
max_throughput_mbps	integer	Maximum throughput defined by this policy. It is specified in terms of Mbps. 0 means no maximum throughput is enforced.

Name	Type	Description
min_throughput_iops	integer	Minimum throughput defined by this policy. It is specified in terms of IOPS. 0 means no minimum throughput is enforced. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

qos_policy

Name	Type	Description
_links	_links	
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.
svm	svm	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	_links	
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.

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	href	

adaptive

Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

Name	Type	Description
absolute_min_iops	integer	Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
expected_iops	integer	Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
peak_iops	integer	Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size.

fixed

QoS policy-groups define a fixed service level objective (SLO) for a storage object.

Name	Type	Description
capacity_shared	boolean	Specifies whether the capacities are shared across all objects that use this QoS policy-group. Default is false.
max_throughput_iops	integer	Maximum throughput defined by this policy. It is specified in terms of IOPS. 0 means no maximum throughput is enforced.
max_throughput_mbps	integer	Maximum throughput defined by this policy. It is specified in terms of Mbps. 0 means no maximum throughput is enforced.
min_throughput_iops	integer	Minimum throughput defined by this policy. It is specified in terms of IOPS. 0 means no minimum throughput is enforced. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

qos_policy

Name	Type	Description
_links	_links	
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

Name	Type	Description
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a QoS policy

GET `/storage/qos/policies/{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
_links	_links	
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	

Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "extreme",
  "object_count": 0,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

adaptive

Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

Name	Type	Description
absolute_min_iops	integer	Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
expected_iops	integer	Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
peak_iops	integer	Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size.

fixed

QoS policy-groups define a fixed service level objective (SLO) for a storage object.

Name	Type	Description
capacity_shared	boolean	Specifies whether the capacities are shared across all objects that use this QoS policy-group. Default is false.
max_throughput_iops	integer	Maximum throughput defined by this policy. It is specified in terms of IOPS. 0 means no maximum throughput is enforced.
max_throughput_mbps	integer	Maximum throughput defined by this policy. It is specified in terms of Mbps. 0 means no maximum throughput is enforced.
min_throughput_iops	integer	Minimum throughput defined by this policy. It is specified in terms of IOPS. 0 means no minimum throughput is enforced. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Update a QoS policy

PATCH `/storage/qos/policies/{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	_links	
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.

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	href	

adaptive

Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

Name	Type	Description
absolute_min_iops	integer	Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
expected_iops	integer	Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
peak_iops	integer	Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size.

fixed

QoS policy-groups define a fixed service level objective (SLO) for a storage object.

Name	Type	Description
capacity_shared	boolean	Specifies whether the capacities are shared across all objects that use this QoS policy-group. Default is false.
max_throughput_iops	integer	Maximum throughput defined by this policy. It is specified in terms of IOPS. 0 means no maximum throughput is enforced.
max_throughput_mbps	integer	Maximum throughput defined by this policy. It is specified in terms of Mbps. 0 means no maximum throughput is enforced.
min_throughput_iops	integer	Minimum throughput defined by this policy. It is specified in terms of IOPS. 0 means no minimum throughput is enforced. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

qos_policy

Name	Type	Description
_links	_links	
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

Name	Type	Description
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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.

– POST /api/storage/qtrees

– GET /api/storage/qtrees

– GET /api/storage/qtrees/{volume-uuid}/{qtree-id}

– PATCH /api/storage/qtrees/{volume-uuid}/{qtree-id}

– DELETE /api/storage/qtrees/{volume-uuid}/{qtree-id}

Examples

Creating a qtree inside a volume for an SVM

This API is used to create a qtree inside a volume for an SVM.

The following example shows how to create a qtree in a FlexVol 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=svm1&volume.name=fv" -H 'accept:
application/hal+json'

# The response
{
  "records": [
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "id": 0,
      "name": "",
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-
005056a7f717/0"
        }
      }
    },
    {
      "svm": {
```

```
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "id": 1,
      "name": "qt1",
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-
005056a7f717/1"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
```

```
      },
      "id": 2,
      "name": "qt2",
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-
005056a7f717/2"
        }
      }
    }
  ],
  "num_records": 3,
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/?svm.name=svm1&volume.name=fv"
    }
  }
}
```

Retrieving properties of a specific qtree using a qtree identifier

This API is used to retrieve properties of a specific qtree using qtree.id.

The following example shows how to use the qtree identifier to retrieve all properties of the qtree using the `fields` query parameter.

```
# The API:
GET /api/storage/qtrees/{volume.uuid}/{id}

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-
005056a7f717/2?fields=*' -H 'accept: application/hal+json'
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
```

```
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
      }
    }
  },
  "id": 2,
  "name": "qt2",
  "security_style": "unix",
  "unix_permissions": 744,
  "export_policy": {
    "name": "default",
    "id": 12884901889,
    "_links": {
      "self": {
        "href": "/api/protocols/nfs/export-policies/12884901889"
      }
    }
  },
  "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": "exp1"
}
}
```

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	_links	

Name	Type	Description
num_records	integer	Number of records
records	array[qtree]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "export_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": 100,
        "name": "default"
      },
      "id": 1,
      "name": "string",
      "path": "/volume3/qtree1",
      "security_style": "string",
      "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	href	
self	href	

_links

Name	Type	Description
self	href	

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

svm

Required in POST

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Required in POST

Name	Type	Description
_links	_links	

Name	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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

qtree

A qtree is a directory at the top level of a volume to which a custom export policy (for fine-grained access control) and a quota rule can be applied, if required.

Name	Type	Description
_links	_links	
export_policy	export_policy	Export Policy
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.
security_style	string	Security style. Valid in POST or PATCH.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	volume	Required in POST

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Create a qtree in a FlexVol or FlexGroup volume

POST `/storage/qtrees`

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	_links	
export_policy	export_policy	Export Policy
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.
security_style	string	Security style. Valid in POST or PATCH.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	volume	Required in POST

Example request


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "id": 1,
  "name": "string",
  "path": "/volume3/qtree1",
  "security_style": "string",
  "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 `svm.uuid` and `svm.name` do not refer to the same SVM.
918236	The specified `volume.uuid` and `volume.name` refer to different volumes.
2621707	No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.
918232	Either `volume.name` or `volume.uuid` must be provided.
5242953	Qtree name must be provided.
5242952	Export Policy ID specified is invalid.

Error Code	Description
5242951	Export Policy supplied does not belong to the specified Export Policy ID.
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	href	

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

svm

Required in POST

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Required in POST

Name	Type	Description
_links	_links	
name	string	The name of the volume.

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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

qtree

A qtree is a directory at the top level of a volume to which a custom export policy (for fine-grained access control) and a quota rule can be applied, if required.

Name	Type	Description
_links	_links	
export_policy	export_policy	Export Policy
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.
security_style	string	Security style. Valid in POST or PATCH.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	volume	Required in POST

job_link

Name	Type	Description
_links	_links	

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve qtree properties

GET `/storage/qtrees/{volume.uuid}/{id}`

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
volume.uuid	string	path	True	Volume UUID
id	string	path	True	Qtree ID
fields	array[string]	query	False	Specify the fields to return.

Response

```
Status: 200, Ok
```

Name	Type	Description
_links	_links	
export_policy	export_policy	Export Policy
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.
security_style	string	Security style. Valid in POST or PATCH.
svm	svm	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,
  "name": "string",
  "path": "/volume3/qtree1",
  "security_style": "string",
  "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	href	

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

svm

Required in POST

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Required in POST

Name	Type	Description
_links	_links	
name	string	The name of the volume.

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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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	_links	
export_policy	export_policy	Export Policy
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.
security_style	string	Security style. Valid in POST or PATCH.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	volume	Required in POST

Example request


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "id": 1,
  "name": "string",
  "path": "/volume3/qtree1",
  "security_style": "string",
  "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	href	

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

svm

Required in POST

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Required in POST

Name	Type	Description
_links	_links	
name	string	The name of the volume.

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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

qtree

A qtree is a directory at the top level of a volume to which a custom export policy (for fine-grained access control) and a quota rule can be applied, if required.

Name	Type	Description
_links	_links	
export_policy	export_policy	Export Policy
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.
security_style	string	Security style. Valid in POST or PATCH.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	volume	Required in POST

job_link

Name	Type	Description
_links	_links	

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Storage quota

Storage quota endpoint overview

Overview

Quotas provide a way to restrict or track the files and space usage by a user, group, or qtree. Quotas are enabled for a specific FlexVol or a FlexGroup volume.

Quotas can have soft or hard limits. Soft limits cause ONTAP to send a notification when specified limits are exceeded. Hard limits prevent a write operation from succeeding when specified limits are exceeded.

Quota policy rule APIs

Quotas are defined as quota policy rules specific to FlexVol or FlexGroup volumes. Each quota rule has a type. The type can be either "user", "group", or "tree".

The following APIs can be used to perform create, retrieve, modify, and delete operations related to quota policy rules for a FlexVol or a FlexGroup volume.

– POST /api/storage/quota/rules

– GET /api/storage/quota/rules

– GET /api/storage/quota/rules/{rule-uuid}

– PATCH /api/storage/quota/rules/{rule-uuid}

– DELETE /api/storage/quota/rules/{rule-uuid}

Enabling and disabling quotas

After the quota rules are created, the quota rules need to be enabled on each FlexVol or FlexGroup volume for soft or hard limits to take effect in the filesystem. Enabling quotas can be done on a volume-by-volume basis.

The following APIs can be used to enable and disable and obtain the quota state for a FlexVol or a FlexGroup volume.

– PATCH /api/storage/volumes/{volume-uuid} -d '{"quota.enabled":"true"}'

– PATCH /api/storage/volumes/{volume-uuid} -d '{"quota.enabled":"false"}'

– GET /api/storage/volumes/{volume-uuid}/?fields=quota.state

Quota report APIs

Quota report records provide usage information for a user, group, or qtree against the quota limits configured on a FlexVol or a FlexGroup volume. The following APIs can be used to retrieve quota reports associated with a FlexVol or a FlexGroup volume.

– GET /api/storage/quota/reports

– GET /api/storage/quota/reports/{volume-uuid}/{index}

Quota resize

Quota resize allows you to modify the quota limits directly in the filesystem.

It is important to note that quota must be enabled on a FlexVol or a FlexGroup volume for `quota resize` to take effect. You can perform a `quota resize` using the quota policy rule PATCH API. If the quota is disabled on the volume, the quota policy rule PATCH API modifies the rule, but this does not affect the limits in the filesystem. The following API can be used to perform `quota resize` provided quota is enabled on the volume.

– PATCH /api/storage/quota/rules/{rule-uuid} The changed limits in the filesystem can be confirmed using the quota report REST API.

– GET /api/storage/quota/reports

Manage storage quota reports

Storage quota reports endpoint overview

Overview

Quota reports provide the current file and space consumption for a user, group, or qtree in a FlexVol or a FlexGroup volume.

Quota report APIs

The following APIs can be used to retrieve quota reports associated with a volume in ONTAP.

– GET /api/storage/quota/reports

– GET /api/storage/quota/reports/{volume_uuid}/{index}

Examples

Retrieving all the quota report records

This API is used to retrieve all the quota report records.

The following example shows how to retrieve quota report records for all 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": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",
        "name": "fg",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
```

```
          }
        }
      },
      "index": 0,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-
005056a7f717/0"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",
        "name": "fg",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
          }
        }
      },
      "index": 1152921504606846976,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-
005056a7f717/1152921504606846976"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
```

```
          }
        }
      },
      "volume": {
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",
        "name": "fg",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
          }
        }
      },
      "index": 3458764513820540928,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-
005056a7f717/3458764513820540928"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",
        "name": "fg",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
          }
        }
      },
      "index": 4611686018427387904,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-
005056a7f717/4611686018427387904"
```

```
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",
        "name": "fg",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
          }
        }
      },
      "index": 5764607523034234880,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-
005056a7f717/5764607523034234880"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "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": 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": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 1153202979583557632,
      "_links": {
        "self": {
```

```
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/1153202979583557632"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 2305843013508661248,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/2305843013508661248"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
```

```
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 3458764513820540928,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/3458764513820540928"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "index": 3459045988797251584,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/3459045988797251584"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
```

```
        "name": "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": "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": 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/34590459887972515 84"  -H 'accept:
application/hal+json'

# Response for quota report group record:
{
"svm": {
  "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
    }
  }
},
"volume": {
  "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
```

```
  "name": "fv",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
    }
  }
},
"index": 3459045988797251584,
"type": "group",
"group": {
  "name"  : "test_group",
  "id"    : "500009"
},
"qtree": {
  "name": "qt1",
  "id": 1,
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
    }
  }
},
"space": {
  "hard_limit": 41943040,
  "soft_limit": 31457280,
  "used": {
    "total": 10567680,
    "soft_limit_percent": 34,
    "hard_limit_percent": 25
  }
}
"files": {
  "soft_limit": 30,
  "hard_limit": 40,
  "used": {
    "total": 11,
    "soft_limit_percent": 37,
    "hard_limit_percent": 28
  }
}
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-
005056a7657c/3459045988797251584"
    }
  }
```

```
}
}
```

Retrieving a single quota report tree record

```
# The 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

Name	Type	In	Required	Description
qtree.id	integer	query	False	Filter by qtree.id
index	integer	query	False	Filter by index
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

Name	Type	In	Required	Description
volume.name	string	query	False	Filter by volume.name
volume.uuid	string	query	False	Filter by volume.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
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
_links	_links	
num_records	integer	Number of records
records	array[quota_report]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "files": {
        "hard_limit": 0,
        "soft_limit": 0,
        "used": {
          "hard_limit_percent": "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": "string",
      "users": [
        {
          "id": "string",
          "name": "string"
        }
      ],
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

used

Name	Type	Description
hard_limit_percent	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	used	

group

This parameter specifies the target group associated with the given quota report record. This parameter is available for group quota records and is not available for user or tree quota records. The target group is identified by a UNIX group name and UNIX group identifer.

Name	Type	Description
id	string	Quota target group ID
name	string	Quota target group name

qtree

This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.

Name	Type	Description
_links	_links	
id	integer	The unique identifier for a qtree.
name	string	The name of the qtree.

used

Name	Type	Description
hard_limit_percent	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	used	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

users

Name	Type	Description
id	string	Quota target user ID
name	string	Quota target user name

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

quota_report

Name	Type	Description
_links	_links	
files	files	
group	group	This parameter specifies the target group associated with the given quota report record. This parameter is available for group quota records and is not available for user or tree quota records. The target group is identified by a UNIX group name and UNIX group identifer.

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 identifer can be a UNIX user identifier or a Windows security identifier.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a specific quota report record

GET `/storage/quota/reports/{volume.uuid}/{index}`

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	_links	
files	files	
group	group	This parameter specifies the target group associated with the given quota report record. This parameter is available for group quota records and is not available for user or tree quota records. The target group is identified by a UNIX group name and UNIX group identifer.
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.

Name	Type	Description
users	array[users]	This parameter specifies the target user or users associated with the given quota report record. This parameter is available for user quota records and is not available for group or tree quota records. The target user or users are identified by a user name and user identifier. The user name can be a UNIX user name or a Windows user name, and the identifer can be a UNIX user identifier or a Windows security identifier.
volume	volume	

Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "files": {
    "hard_limit": 0,
    "soft_limit": 0,
    "used": {
      "hard_limit_percent": "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": "string",
  "users": [
    {
      "id": "string",
      "name": "string"
    }
  ],
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```


Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
918235	A volume with UUID was not found.

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

used

Name	Type	Description
hard_limit_percent	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	used	

group

This parameter specifies the target group associated with the given quota report record. This parameter is available for group quota records and is not available for user or tree quota records. The target group is identified by a UNIX group name and UNIX group identifer.

Name	Type	Description
id	string	Quota target group ID
name	string	Quota target group name

qtree

This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a

user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.

Name	Type	Description
_links	_links	
id	integer	The unique identifier for a qtree.
name	string	The name of the qtree.

used

Name	Type	Description
hard_limit_percent	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	used	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

users

Name	Type	Description
id	string	Quota target user ID
name	string	Quota target user name

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage storage quota policy rules

Storage quota rules endpoint overview

Overview

Quotas are defined in quota rules specific to FlexVols or FlexGroup volumes. Each quota rule has a type. The type can be "user", "group", or "tree".

– User rules must have the user property and qtree property.

– Group rules must have the group property and qtree property.

– Tree rules must have the qtree property and not have the user or group property.

Quota policy rule APIs

The following APIs can be used to perform create, retrieve, modify, and delete operations related to quota policy rules.

– POST /api/storage/quota/rules

– GET /api/storage/quota/rules

– GET /api/storage/quota/rules/{rule-uuid}

– PATCH /api/storage/quota/rules/{rule-uuid}

– DELETE /api/storage/quota/rules/{rule-uuid}

Examples

Retrieving all quota policy rules

This API is used to retrieve all quota policy rules.

The following example shows how to retrieve quota policy rules for 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": "svm1",
```

```
        "_links": {
          "self": {
            "href": "/api/svm/svms/038545f8-9ff8-11e8-bce6-005056a73bed"
          }
        }
      },
      "volume": {
        "uuid": "ab3df793-0f02-43c6-9514-4f142fc8cc92",
        "name": "vol1",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/ab3df793-0f02-43c6-9514-
4f142fc8cc92"
          }
        }
      },
      "uuid": "66319cbe-b837-11e8-9c5a-005056a7e88c",
      "_links": {
        "self": {
          "href": "/api/storage/quota/rules/66319cbe-b837-11e8-9c5a-
005056a7e88c"
        }
      }
    },
    {
      "svm": {
        "uuid": "038545f8-9ff8-11e8-bce6-005056a73bed",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/038545f8-9ff8-11e8-bce6-005056a73bed"
          }
        }
      },
      "volume": {
        "uuid": "ab3df793-0f02-43c6-9514-4f142fc8cc92",
        "name": "vol1",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/ab3df793-0f02-43c6-9514-
4f142fc8cc92"
          }
        }
      },
      "uuid": "dbd5b443-b7a4-11e8-bc58-005056a7e88c",
      "_links": {
```

```
        "self": {
          "href": "/api/storage/quota/rules/dbd5b443-b7a4-11e8-bc58-
005056a7e88c"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules"
    }
  }
}
```

Retrieving a specific quota policy rule

This API is used to retrieve a quota policy rule for a specific qtree.

The following example shows how to retrieve a quota policy user rule for a specific qtree.

```
# The API:
GET /api/storage/quota/rules/{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": "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": "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": "vol1",
    "_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"
    }
  }
}
```

```
# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/49b1134f-19ab-11e9-
a04d-005056a72f42' -H 'accept: application/hal+json'

# Response for a tree rule for a specific qtree:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-
ab489830d268"
      }
    }
  },
  "uuid": "49b1134f-19ab-11e9-a04d-005056a72f42",
  "type": "tree",
  "qtree": {
    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/c1b64eea-ca8b-45ec-9397-
ab489830d268/1"
      }
    }
  },
  "space": {
    "hard_limit": 1048576,
    "soft_limit": 838861
  },
  "files": {
    "hard_limit": 100,
```

```
    "soft_limit": 40
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules/49b1134f-19ab-11e9-a04d-
005056a72f42"
    }
  }
}
```

Retrieving a quota policy group rule for a specific qtree

```
# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/b9236852-19ab-11e9-
a04d-005056a72f42' -H 'accept: application/hal+json'

# Response for a group rule:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-
ab489830d268"
      }
    }
  },
  "uuid": "b9236852-19ab-11e9-a04d-005056a72f42",
  "type": "group",
  "group": {"name" : "group1"},
  "qtree": {
    "name": "qt1",
```

```
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/c1b64eea-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": "vol1"
},
"type": "user",
"users": [ {"name" : "jsmith"} ],
"qtree": {
  "name":"qt1"
},
"user_mapping": "on",
"space": {
  "hard_limit": 8192,
  "soft_limit": 1024
},
"files": {
  "hard_limit": 20,
  "soft_limit": 10
}
}

# The response
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1"
      },
      "volume": {
        "name": "fv"
      },
      "uuid": "3220eea6-5049-11e9-bfb7-005056a7f717",
      "type": "user",
      "users": [
        {
          "name" : "jsmith"
        }
      ],
      "qtree": {
        "name": "qt1"
      },
      "user_mapping": "on",
      "space": {
        "hard_limit": 8192,
        "soft_limit": 1024
      },
      "files": {
        "hard_limit": 20,
```

```
        "soft_limit": 10
      },
      "_links": {
        "self": {
          "href": "/api/storage/quota/rules/3220eea6-5049-11e9-bfb7-
005056a7f717"
        }
      }
    }
  ],
  "job": {
    "uuid": "32223924-5049-11e9-bfb7-005056a7f717",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/32223924-5049-11e9-bfb7-005056a7f717"
      }
    }
  }
}
```

Creating a quota policy group rule using POST.

```
# The API:
POST /api/storage/quota/rules

# The call:
curl -X POST 'https://<mgmt-
ip>/api/storage/quota/rules?return_records=true' -H 'accept:
application/hal+json' -d @test_quota_post.txt
test_quota_post.txt(body):
{
"svm": {
  "name": "svm1"
},
"volume": {
  "name": "vol1"
},
"type": "group",
"group": {
  "name" : "test_group1"
}
"qtree": {
```

```
  "name":"qt1"
},
"space": {
  "hard_limit": 8192,
  "soft_limit": 1024
},
"files": {
  "hard_limit": 20,
  "soft_limit": 10
}
}

# The response
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1"
      },
      "volume": {
        "name": "fv"
      },
      "uuid": "3b130f7d-504a-11e9-bfb7-005056a7f717",
      "type": "group",
      "group": {
          "name" : "test_group1"
      },
      "qtree": {
        "name": "qt1"
      },
      "space": {
        "hard_limit": 8192,
        "soft_limit": 1024
      },
      "files": {
        "hard_limit": 20,
        "soft_limit": 10
      },
      "_links": {
        "self": {
          "href": "/api/storage/quota/rules/3b130f7d-504a-11e9-bfb7-
005056a7f717"
        }
      }
    }
```

```
  ],
  "job": {
    "uuid": "32223924-5049-11e9-bfb7-005056a7f717",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/32223924-5049-11e9-bfb7-005056a7f717"
      }
    }
  }
}
```

Creating a quota policy tree rule using POST

```
# The API:
POST /api/storage/quota/rules

# The call:
curl -X POST 'https://<mgmt-
ip>/api/storage/quota/rules?return_records=true' -H 'accept:
application/hal+json' -d @test_quota_post.txt
test_quota_post.txt(body):
{
"svm": {
  "name": "svm1"
},
"volume": {
  "name": "vol1"
},
"type": "tree",
"qtree": {
  "name":"qt1"
},
"space": {
  "hard_limit": 8192,
  "soft_limit": 1024
},
"files": {
  "hard_limit": 20,
  "soft_limit": 10
}
}
```

```
# The response
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1"
      },
      "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

Name	Type	In	Required	Description
space.soft_limit	integer	query	False	Filter by space.soft_limit
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	_links	
num_records	integer	Number of records
records	array[quota_rule]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "group": {
        "id": "string",
        "name": "string"
      },
      "qtree": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": 1,
        "name": "qt1"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "users": [
        {
          "id": "string",
          "name": "string"
        }
```

```
      ],
      "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

files

Name	Type	Description
hard_limit	integer	This parameter specifies the hard limit for files. This is valid in POST or PATCH.
soft_limit	integer	This parameter specifies the soft limit for files. This is valid in POST or PATCH.

group

This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
id	string	Quota target group ID
name	string	Quota target group name

qtree

This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.

Name	Type	Description
_links	_links	
id	integer	The unique identifier for a qtree.
name	string	The name of the qtree.

space

Name	Type	Description
hard_limit	integer	This parameter specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH.
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	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

users

Name	Type	Description
id	string	Quota target user ID
name	string	Quota target user name

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

quota_rule

Name	Type	Description
_links	_links	
files	files	
group	group	This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
qtree	qtree	This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.
space	space	
svm	svm	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 identifer can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Create a quota policy rule for a FlexVol or a FlexGroup volume

POST `/storage/quota/rules`

Creates a quota policy rule for a FlexVol or a FlexGroup volume.

Important notes:

- Unlike CLI/ONTAPI, the `quota policy` input is not needed for POST.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the qtree.
- `volume.uuid` or `volume.name` - Existing volume in which to create the qtree.
- `type` - Quota type for the rule. This type can be `user`, `group`, or `tree`.
- `users.name` or `user.id` - If the quota type is user, this property takes the user name or user ID. For default user quota rules, the user name must be specified as "".
- `group.name` or `group.id` - If the quota type is group, this property takes the group name or group ID. For default group quota rules, the group name must be specified as "".
- `qtree.name` - Qtree for which to create the rule. For default tree rules, the qtree name must be specified as "".

Recommended optional properties

- `space.hard_limit` - Specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes.
- `space.soft_limit` - Specifies the space soft limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes.
- `files.hard_limit` - Specifies the hard limit for files.
- `files.hard_limit` - Specifies the soft limit for files.
- `user_mapping` - Specifies the user_mapping. This property is valid only for quota policy rules of type `user`.

Related ONTAP commands

- `quota policy rule create`

Learn more

- DOC /storage/quota/rules

Request Body

Name	Type	Description
_links	_links	
files	files	
group	group	This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.
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.

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 identifer 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


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "group": {
    "id": "string",
    "name": "string"
  },
  "qtree": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 1,
    "name": "qt1"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "users": [
    {
      "id": "string",
      "name": "string"
    }
  ],
  "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
```

```
    }
}
```

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 `svm.uuid` and `svm.name` do not refer to the same SVM.
918236	The specified `volume.uuid` and `volume.name` refer to different volumes.
2621707	No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.

Error Code	Description
918232	Either `volume.name` or `volume.uuid` must be provided.
5308549	Volume does not belong to SVM.
5308552	Failed to get default quota policy name for SVM.
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	href	

files

Name	Type	Description
hard_limit	integer	This parameter specifies the hard limit for files. This is valid in POST or PATCH.
soft_limit	integer	This parameter specifies the soft limit for files. This is valid in POST or PATCH.

group

This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
id	string	Quota target group ID
name	string	Quota target group name

qtree

This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.

Name	Type	Description
_links	_links	
id	integer	The unique identifier for a qtree.
name	string	The name of the qtree.

space

Name	Type	Description
hard_limit	integer	This parameter specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH.
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	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

users

Name	Type	Description
id	string	Quota target user ID
name	string	Quota target user name

volume

Name	Type	Description
_links	_links	

Name	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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

quota_rule

Name	Type	Description
_links	_links	
files	files	
group	group	This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
qtree	qtree	This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.
space	space	
svm	svm	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 identifer 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[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a quota policy rule

DELETE `/storage/quota/rules/{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	job_link	

Example response

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

Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
5308545	The specified quota rule UUID is invalid.
5308561	Failed to obtain volume quota state or invalid quota state obtained for volume.
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	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve quota policy rule properties

GET `/storage/quota/rules/{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
_links	_links	
files	files	
group	group	This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
qtree	qtree	This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.
space	space	
svm	svm	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 identifer can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.
volume	volume	

Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "group": {
    "id": "string",
    "name": "string"
  },
  "qtree": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 1,
    "name": "qt1"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "users": [
    {
      "id": "string",
      "name": "string"
    }
  ],
  "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
```

```
    }
}
```

Error

```
Status: Default
```

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	href	

files

Name	Type	Description
hard_limit	integer	This parameter specifies the hard limit for files. This is valid in POST or PATCH.
soft_limit	integer	This parameter specifies the soft limit for files. This is valid in POST or PATCH.

group

This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
id	string	Quota target group ID
name	string	Quota target group name

qtree

This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.

Name	Type	Description
_links	_links	
id	integer	The unique identifier for a qtree.
name	string	The name of the qtree.

space

Name	Type	Description
hard_limit	integer	This parameter specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH.
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	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

users

Name	Type	Description
id	string	Quota target user ID
name	string	Quota target user name

volume

Name	Type	Description
_links	_links	

Name	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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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
_links	_links	
files	files	
group	group	This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
qtree	qtree	This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.
space	space	
svm	svm	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 identifer can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.
volume	volume	

Example 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": "string",
  "users": [
    {
      "id": "string",
      "name": "string"
    }
  ],
  "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
```

```
    }
}
```

Response

```
Status: 202, Accepted
```

Name	Type	Description
job	job_link	

Example response

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

Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
5308545	The specified quota rule UUID is invalid.
5308561	Failed to obtain volume quota state or invalid quota state obtained for volume.
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.

Error Code	Description
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	href	

files

Name	Type	Description
hard_limit	integer	This parameter specifies the hard limit for files. This is valid in POST or PATCH.
soft_limit	integer	This parameter specifies the soft limit for files. This is valid in POST or PATCH.

group

This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
id	string	Quota target group ID
name	string	Quota target group name

qtree

This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.

Name	Type	Description
_links	_links	
id	integer	The unique identifier for a qtree.
name	string	The name of the qtree.

space

Name	Type	Description
hard_limit	integer	This parameter specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH.
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	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

users

Name	Type	Description
id	string	Quota target user ID
name	string	Quota target user name

volume

Name	Type	Description
_links	_links	

Name	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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

quota_rule

Name	Type	Description
_links	_links	
files	files	
group	group	This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
qtree	qtree	This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.
space	space	
svm	svm	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 identifer 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[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage shelves

Storage shelves endpoint overview

Retrieving storage shelf information

The storage shelf GET API retrieves all of the shelves in the cluster.

Examples

1) Retrieve a list of shelves from the cluster

The following example shows the response with a list of shelves in the cluster:

```
# The API:
/api/storage/shelves

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

# The response:
```

```
{
"records": [
  {
    "uid": "3109174803597886800",
    "_links": {
      "self": {
        "href": "/api/storage/shelves/3109174803597886800"
      }
    }
  },
  {
    "uid": "9237728366621690448",
    "_links": {
      "self": {
        "href": "/api/storage/shelves/9237728366621690448"
      }
    }
  },
  {
    "uid": "9946762738829886800",
    "_links": {
      "self": {
        "href": "/api/storage/shelves/9946762738829886800"
      }
    }
  },
  {
    "uid": "10318311901725526608",
    "_links": {
      "self": {
        "href": "/api/storage/shelves/10318311901725526608"
      }
    }
  },
  {
    "uid": "13477584846688355664",
    "_links": {
      "self": {
        "href": "/api/storage/shelves/13477584846688355664"
      }
    }
  }
],
"num_records": 5,
"_links": {
  "self": {
```

```
    "href": "/api/storage/shelves/"
  }
}
}
```

2) Retrieve a specific shelf from the cluster

The following example shows the response of the requested shelf. If there is no shelf with the requested uid, an error is returned.

```
# The API:
/api/storage/shelves/{uid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/shelves/3109174803597886800" -H
"accept: application/hal+json"

# The response:
{
"uid": "3109174803597886800",
"name": "6.10",
"id": "10",
"serial_number": "SHU0954292N0HAH",
"model": "DS4246",
"module_type": "iom6",
"internal": false,
"state": "ok",
"connection_type": "sas",
"disk_count": 24,
"paths": [
  {
    "name": "0e",
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
      }
    },
    "_links": {
```

```
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0e"
      }
    }
  },
  {
    "name": "0g",
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0g"
      }
    }
  }
],
"bays": [
  {
    "id": 0,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 1,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 2,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
```

```
    "id": 3,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 4,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 5,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 6,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 7,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 8,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 9,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 10,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
```

```
  },
  {
    "id": 11,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 12,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 13,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 14,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 15,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 16,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 17,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 18,
    "has_disk": true,
```

```
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 19,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 20,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 21,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 22,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 23,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    }
  ],
  "frus": [
    {
      "type": "module",
      "id": 0,
      "state": "ok",
      "part_number": "111-00690+B2",
      "serial_number": "8001900099",
      "firmware_version": "0191"
    },
    {
      "type": "module",
      "id": 1,
```

```
      "state": "ok",
      "part_number": "111-00190+B0",
      "serial_number": "7903785183",
      "firmware_version": "0191"
    },
    {
      "type": "psu",
      "id": 1,
      "state": "ok",
      "part_number": "0082562-12",
      "serial_number": "PMW82562007513E",
      "firmware_version": "0311"
    },
    {
      "type": "psu",
      "id": 2,
      "state": "ok",
      "part_number": "0082562-12",
      "serial_number": "PMW825620075138",
      "firmware_version": "0311"
    },
    {
      "type": "psu",
      "id": 3,
      "state": "ok",
      "part_number": "0082562-12",
      "serial_number": "PMW8256200750BA",
      "firmware_version": "0311"
    },
    {
      "type": "psu",
      "id": 4,
      "state": "ok",
      "part_number": "0082562-12",
      "serial_number": "PMW8256200750A2",
      "firmware_version": "0311"
    }
  ],
  "ports": [
    {
      "id": 0,
      "module_id": "a",
      "designator": "square",
      "state": "connected",
      "internal": false,
      "wwn": "500A098000C9EDBF",
```

```
    "cable": {
      "identifier": "5001086000702488-500a098000c9edbf",
      "part_number": "112-00430+A0",
      "length": "2m",
      "serial_number": "APF16510229807"
    },
    "remote": {
      "wwn": "5001086000702488",
      "phy": "08"
    }
  },
  {
    "id": 1,
    "module_id": "a",
    "designator": "circle",
    "state": "connected",
    "internal": false,
    "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	_links	
num_records	integer	Number of records
records	array[shelf]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "bays": [
        {
          "id": 0,
          "state": "ok",
          "type": "single_disk"
        }
      ],
      "connection_type": "sas",
      "disk_count": 12,
      "drawers": [
        {
          "disk_count": 12,
          "error": "string",
          "part_number": "111-03071",
          "serial_number": "021604008263",
          "state": "ok"
        }
      ],
      "frus": [
        {
          "firmware_version": "0191",
          "part_number": "111-00690+A2",
          "serial_number": 8000166294,
          "state": "error",
          "type": "module"
        }
      ],
      "id": 1,
      "model": "DS2246",
      "module_type": "iom6",
      "name": 1.1,
      "paths": [
        {
```

```
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "2a",
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "ports": [
        {
          "cable": {
            "identifier": "500a0980000b6c3f-50000d1703544b80",
            "length": "2m",
            "part_number": "112-00431+A0",
            "serial_number": 616930439
          },
          "designator": "square",
          "id": 0,
          "mac_address": "string",
          "module_id": "a",
          "remote": {
            "chassis": "string",
            "mac_address": "string",
            "phy": 12,
            "port": "string",
            "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	href	
self	href	

bays

Name	Type	Description
has_disk	boolean	
id	integer	
state	string	
type	string	

drawers

Name	Type	Description
closed	boolean	
disk_count	integer	
error	string	
id	integer	
part_number	string	
serial_number	string	
state	string	

frus

Name	Type	Description
firmware_version	string	
id	integer	
part_number	string	
serial_number	string	
state	string	

Name	Type	Description
type	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

paths

Storage port

Name	Type	Description
_links	_links	
name	string	
node	node	

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

remote

Name	Type	Description
chassis	string	
mac_address	string	
phy	string	
port	string	
wwn	string	

ports

Name	Type	Description
cable	cable	
designator	string	
id	integer	
internal	boolean	
mac_address	string	
module_id	string	
remote	remote	
state	string	
wwn	string	

shelf

Name	Type	Description
bays	array[bays]	
connection_type	string	
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	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a shelf

GET `/storage/shelves/{uid}`

Retrieves a specific shelf.

Related ONTAP commands

- `storage shelf show`
- `storage shelf port show`
- `storage shelf drawer show`

Learn more

- DOC /storage/shelves

Parameters

Name	Type	In	Required	Description
uid	string	path	True	Shelf UID
fields	array[string]	query	False	Specify the fields to return.

Response

```
Status: 200, Ok
```

Name	Type	Description
bays	array[bays]	
connection_type	string	

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,
      "error": "string",
      "part_number": "111-03071",
      "serial_number": "021604008263",
      "state": "ok"
    }
  ],
  "frus": [
    {
      "firmware_version": "0191",
      "part_number": "111-00690+A2",
      "serial_number": 8000166294,
      "state": "error",
      "type": "module"
    }
  ],
  "id": 1,
  "model": "DS2246",
  "module_type": "iom6",
  "name": 1.1,
  "paths": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "2a",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
```

```
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ],
    "ports": [
      {
        "cable": {
          "identifier": "500a0980000b6c3f-50000d1703544b80",
          "length": "2m",
          "part_number": "112-00431+A0",
          "serial_number": 616930439
        },
        "designator": "square",
        "id": 0,
        "mac_address": "string",
        "module_id": "a",
        "remote": {
          "chassis": "string",
          "mac_address": "string",
          "phy": 12,
          "port": "string",
          "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	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

paths

Storage port

Name	Type	Description
_links	_links	
name	string	
node	node	

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

remote

Name	Type	Description
chassis	string	
mac_address	string	
phy	string	
port	string	
wwn	string	

ports

Name	Type	Description
cable	cable	
designator	string	
id	integer	
internal	boolean	
mac_address	string	

Name	Type	Description
module_id	string	
remote	remote	
state	string	
wwn	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage Snapshot copy policies

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:

– POST /api/storage/snapshot_policies

– GET /api/storage/snapshot_policies

– GET /api/storage/snapshot_policies/{uuid}

– PATCH /api/storage/snapshot_policies/{uuid}

– DELETE /api/storage/snapshot_policies/{uuid}

Examples

Creating a Snapshot copy policy

The POST operation is used to create a Snapshot copy policy with the specified attributes.

```
# The API:
/api/storage/snapshot_policies

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/snapshot_policies" -H
"accept: application/hal+json" -d '{"name": "new_policy", "enabled":
"true", "comment": "policy comment", "copies": [{ "schedule": { "name":
"5min" }, "count": "5", "prefix": "xyz" }], "svm": { "name": "vs0"}}'

# The response:
HTTP/1.1 201 Created
Date: Tue, 12 Mar 2019 21:20:24 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/snapshot_policies/a69d8173-450c-11e9-aa44-
005056bbc848
Content-Length: 369
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "uuid": "a69d8173-450c-11e9-aa44-005056bbc848",
    "svm": {
      "name": "vs0"
    },
    "name": "new_policy",
    "comment": "This is a 5min schedule policy",
    "enabled": true,
    "copies": [
      {
        "count": 5,
        "schedule": {
          "name": "5min"
        }
      }
    ]
  }
]
}
```

Retrieving Snapshot copy policy attributes

The GET operation is used to retrieve Snapshot copy policy attributes.

```
# The API:
/api/storage/snapshot_policies

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

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:17:17 GMT
Server: libzapid-http
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 686
Content-Type: application/json
{
"records": [
  {
    "uuid": "0fa7a554-348d-11e9-b55e-005056bbf1c8",
    "name": "spsv0",
    "_links": {
      "self": {
        "href": "/api/storage/snapshot_policies/0fa7a554-348d-11e9-b55e-
005056bbf1c8"
      }
    }
  },
  {
    "uuid": "3c112527-2fe8-11e9-b55e-005056bbf1c8",
    "name": "default",
    "_links": {
      "self": {
        "href": "/api/storage/snapshot_policies/3c112527-2fe8-11e9-b55e-
005056bbf1c8"
      }
    }
  },
  {
    "uuid": "3c1c1656-2fe8-11e9-b55e-005056bbf1c8",
    "name": "default-1weekly",
    "_links": {
      "self": {
        "href": "/api/storage/snapshot_policies/3c1c1656-2fe8-11e9-b55e-
005056bbf1c8"
      }
    }
```

```
  },
  {
    "uuid": "3c228b82-2fe8-11e9-b55e-005056bbf1c8",
    "name": "none",
    "_links": {
      "self": {
        "href": "/api/storage/snapshot_policies/3c228b82-2fe8-11e9-b55e-
005056bbf1c8"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/storage/snapshot_policies/"
  }
}
}
```

Retrieving the attributes of a specific Snapshot copy policy

The GET operation is used to retrieve the attributes of a specific Snapshot copy policy.

```
# The API:
/api/storage/snapshot_policies

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/snapshot_policies/3c112527-
2fe8-11e9-b55e-005056bbf1c8" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:24:48 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 381
Content-Type: application/json
{
"uuid": "3c112527-2fe8-11e9-b55e-005056bbf1c8",
"name": "default",
"comment": "Default policy with hourly, daily & weekly schedules.",
"enabled": true,
"scope": "cluster",
"copies": [
```

```
  {
    "count": 6,
    "prefix": "hourly",
    "schedule": {
      "name": "hourly"
    }
  },
  {
    "count": 2,
    "prefix": "daily",
    "schedule": {
      "name": "daily"
    }
  },
  {
    "count": 2,
    "prefix": "weekly",
    "schedule": {
      "name": "weekly"
    }
  }
],
"_links": {
  "self": {
    "href": "/api/storage/snapshot_policies/3c112527-2fe8-11e9-b55e-
005056bbf1c8"
  }
}
}
```

Updating a Snapshot copy policy

The PATCH operation is used to update the specific attributes of a Snapshot copy policy.

```
# The API:
/api/storage/snapshot_policies/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/snapshot_policies/ae9e65c4-
4506-11e9-aa44-005056bbc848" -d '{"enabled": "false" }' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:27:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
```

Deleting a Snapshot copy policy

The DELETE operation is used to delete a Snapshot copy policy.

```
# The API:
/api/storage/snapshot_policies/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/snapshot_policies/
ae9e65c4-4506-11e9-aa44-005056bbc848" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:19:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
```

Retrieve Snapshot copy policies

GET `/storage/snapshot-policies`

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	_links	
num_records	integer	Number of records
records	array[snapshot_policy]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "copies": [
        {
          "prefix": "string",
          "schedule": {
            "name": "hourly"
          },
          "snapmirror_label": "string"
        }
      ],
      "enabled": 1,
      "name": "default",
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

schedule

Name	Type	Description
name	string	Schedule at which Snapshot copies are captured on the volume. Some common schedules already defined in the system are hourly, daily, weekly, at 15 minute intervals, and at 5 minute intervals. Snapshot copy policies with custom schedules can be referenced.

copies

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapshot_policy

The Snapshot copy policy object is associated with a read-write volume used to create and delete Snapshot copies at regular intervals.

Name	Type	Description
_links	_links	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[copies]	
enabled	boolean	Is the Snapshot copy policy enabled?
name	string	Name of the Snapshot copy policy.
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
svm	svm	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[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a Snapshot copy policy

DELETE `/storage/snapshot-policies/{uuid}`

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.

Error Code	Description
1638416	Cannot delete policy. Reason: Cannot verify whether policy is in use.
1638430	Cannot delete policy. Reason: Policy is in use by at least one Vserver.
1638430	Cannot delete built-in policy.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve Snapshot copy policy details

GET `/storage/snapshot-policies/{uuid}`

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	_links	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[copies]	
enabled	boolean	Is the Snapshot copy policy enabled?
name	string	Name of the Snapshot copy policy.
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	

Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "copies": [
    {
      "prefix": "string",
      "schedule": {
        "name": "hourly"
      },
      "snapmirror_label": "string"
    }
  ],
  "enabled": 1,
  "name": "default",
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

schedule

Name	Type	Description
name	string	Schedule at which Snapshot copies are captured on the volume. Some common schedules already defined in the system are hourly, daily, weekly, at 15 minute intervals, and at 5 minute intervals. Snapshot copy policies with custom schedules can be referenced.

copies

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.

Name	Type	Description
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Update a Snapshot copy policy

PATCH `/storage/snapshot-policies/{uuid}`

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
_links	_links	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[copies]	
enabled	boolean	Is the Snapshot copy policy enabled?
name	string	Name of the Snapshot copy policy.
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	

Example request


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "copies": [
    {
      "prefix": "string",
      "schedule": {
        "name": "hourly"
      },
      "snapmirror_label": "string"
    }
  ],
  "enabled": 1,
  "name": "default",
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


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	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

schedule

Name	Type	Description
name	string	Schedule at which Snapshot copies are captured on the volume. Some common schedules already defined in the system are hourly, daily, weekly, at 15 minute intervals, and at 5 minute intervals. Snapshot copy policies with custom schedules can be referenced.

copies

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
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	_links	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[copies]	
enabled	boolean	Is the Snapshot copy policy enabled?
name	string	Name of the Snapshot copy policy.
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
svm	svm	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[error_arguments]	Message arguments
code	string	Error code

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

Create a Snapshot copy policy

POST `/storage/snapshot-policies`

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
_links	_links	

Name	Type	Description
comment	string	A comment associated with the Snapshot copy policy.
copies	array[copies]	
enabled	boolean	Is the Snapshot copy policy enabled?
name	string	Name of the Snapshot copy policy.
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	

Example request


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "copies": [
    {
      "prefix": "string",
      "schedule": {
        "name": "hourly"
      },
      "snapmirror_label": "string"
    }
  ],
  "enabled": 1,
  "name": "default",
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


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	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

schedule

Name	Type	Description
name	string	Schedule at which Snapshot copies are captured on the volume. Some common schedules already defined in the system are hourly, daily, weekly, at 15 minute intervals, and at 5 minute intervals. Snapshot copy policies with custom schedules can be referenced.

copies

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
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	_links	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[copies]	
enabled	boolean	Is the Snapshot copy policy enabled?
name	string	Name of the Snapshot copy policy.
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
svm	svm	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[error_arguments]	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:

– PATCH /api/storage/volumes/{volume-uuid} -d '{"quota.enabled":"true"}'

– PATCH /api/storage/volumes/{volume-uuid} -d '{"quota.enabled":"false"}'

– GET /api/storage/volumes/{volume-uuid}/?fields=quota.state

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:

– POST /api/storage/volumes

– GET /api/storage/volumes

– GET /api/storage/volumes/{uuid}

– PATCH /api/storage/volumes/{uuid}

– DELETE /api/storage/volumes/{uuid}

Examples

Creating a volume

The POST request is used to create a new volume and to specify its properties.

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

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name" : "vol1", "aggregates.name" : ["aggr1"],
"svm.name" : "vs1"}'

# The response:
{
"job": {
  "uuid": "b89bc5dd-94a3-11e8-a7a3-0050568edf84",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/b89bc5dd-94a3-11e8-a7a3-0050568edf84"
    }
  }
}
}
```

Creating a 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" : "vol1", "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":"vol1_clone",{"clone":"parent_volume":
{"name": "vol1"}},"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=vol1_clone
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "c9ee0040-39a5-11e9-9b24-00a098439a83",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c9ee0040-39a5-11e9-9b24-00a098439a83"
    }
  }
}
}
```

Volumes reported in the GET REST API

The following types of volumes are reported:

– RW, DP and LS volume

– FlexGroup volume

– FlexCache volume

– FlexClone volume

The following types of volumes are not reported:

– DEL volume

– TEMP volume

– Node Root volume

– System Vserver volume

– FlexGroup constituent

– FlexCache constituent

Examples

Retrieving the 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

Name	Type	In	Required	Description
error_state.is_inconsistent	boolean	query	False	Filter by error_state.is_inconsistent
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

Name	Type	In	Required	Description
snaplock.compliance_clock_time	string	query	False	Filter by snaplock.compliance_clock_time
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

Name	Type	In	Required	Description
metric.throughput.total	integer	query	False	Filter by metric.throughput.total
metric.throughput.other	integer	query	False	Filter by metric.throughput.other
metric.throughput.read	integer	query	False	Filter by metric.throughput.read
metric.throughput.write	integer	query	False	Filter by metric.throughput.write
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_complete	string	query	False	Filter by movement.percent_complete
movement.destination_aggregate.uuid	string	query	False	Filter by movement.destination_aggregate.uuid
movement.destination_aggregate.name	string	query	False	Filter by movement.destination_aggregate.name
movement.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

Name	Type	In	Required	Description
efficiency.cross_volume_dedupe	string	query	False	Filter by efficiency.cross_volume_dedupe
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

Name	Type	In	Required	Description
space.logical_space.available	integer	query	False	Filter by space.logical_space.available
space.logical_space.reporting	boolean	query	False	Filter by space.logical_space.reporting
space.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

Name	Type	In	Required	Description
statistics.iops_raw.total	integer	query	False	Filter by statistics.iops_raw.total
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

Name	Type	In	Required	Description
statistics.latency_raw.read	integer	query	False	Filter by statistics.latency_raw.read
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

Name	Type	In	Required	Description
clone.parent_snapshot.uuid	string	query	False	Filter by clone.parent_snapshot.uuid
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	Type	In	Required	Description
name	string	query	False	Filter by name
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.

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	_links	
num_records	integer	Number of records
records	array[volume]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "aggregates": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "application": {
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
      },
      "autosize": {
        "mode": "string"
      },
      "clone": {
        "parent_snapshot": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "this_snapshot",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
```

```
      "parent_svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "parent_volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      },
      "split_complete_percent": 0,
      "split_estimate": 0
    },
    "comment": "string",
    "create_time": "2018-06-04 19:00:00 UTC",
    "efficiency": {
      "compaction": "string",
      "compression": "string",
      "cross_volume_dedupe": "string",
      "dedupe": "string"
    },
    "encryption": {
      "key_id": "string",
      "state": "string",
      "status": {
        "code": "string",
        "message": "string"
      },
      "type": "string"
    },
    "files": {
      "used": 0
    },
    "flexcache_endpoint_type": "string",
    "guarantee": {
      "type": "string"
    },
    "language": "string",
```

```
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "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": "string"
      },
      "name": "vol_cs_dept",
      "nas": {
        "export_policy": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
```

```
        }
      },
      "id": 100,
      "name": "default"
    },
    "path": "/user/my_volume",
    "security_style": "string",
    "unix_permissions": 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": "string"
  },
  "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": "string",
    "statistics": {
      "iops_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 11:20:13 UTC"
    },
    "style": "string",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
```

```
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tiering": {
        "policy": "string"
      },
      "type": "string",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ]
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

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 threhold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size..
maximum	integer	Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume.
minimum	integer	Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume.
mode	string	Autosize mode for the volume. grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value. grow_shrink ‐ Volume grows or shrinks in response to the amount of space used. off ‐ Autosizing of the volume is disabled.

Name	Type	Description
shrink_threshold	integer	Used space threshold size, in percentage, for the automatic shrinkage of the volume. When the amount of used space in the volume drops below this threshold, the volume automatically shrinks unless it has reached the minimum size. The volume shrinks when the 'space.used' is less than the 'shrink_threshold' percent of 'space.size'. The 'shrink_threshold' size cannot be greater than or equal to the 'grow_threshold' size.

snapshot_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

parent_svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

parent_volume

Name	Type	Description
_links	_links	
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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

clone

Name	Type	Description
is_flexclone	boolean	Specifies if this volume is a normal FlexVol or FlexClone. This field needs to be set when creating a FlexClone. Valid in POST.
parent_snapshot	snapshot_reference	
parent_svm	parent_svm	SVM, applies only to SVM-scoped objects.
parent_volume	parent_volume	
split_complete_percent	integer	Percentage of FlexClone blocks split from its parent volume.
split_estimate	integer	Space required by the containing-aggregate to split the FlexClone volume.
split_initiated	boolean	This field is set when split is executed on any FlexClone, that is when the FlexClone volume is split from its parent FlexVol. This field needs to be set for splitting a FlexClone form FlexVol. Valid in PATCH.

efficiency

Name	Type	Description
compaction	string	The system can be enabled/disabled compaction. inline ‐ Data will be compacted first and written to the volume. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.
compression	string	The system can be enabled/disabled compression. inline ‐ Data will be compressed first and written to the volume. background ‐ Data will be written to the volume and compressed later. both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.
cross_volume_dedupe	string	The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. inline ‐ Data will be cross volume deduped first and written to the volume. background ‐ Data will be written to the volume and cross volume deduped later. both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.

Name	Type	Description
dedupe	string	The system can be enabled/disabled dedupe. inline ‐ Data will be deduped first and written to the volume. background ‐ Data will be written to the volume and deduped later. both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.

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.

Name	Type	Description
state	string	Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one.
status	status	
type	string	Volume encryption type. none ‐ The volume is a plain-text one. volume ‐ The volume is encrypted with volume key (NVE volume). aggregate ‐ 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 ‐ File system is inconsistent. false ‐ 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?
type	string	The type of space guarantee of this volume in the aggregate.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Peformance 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	Peformance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Peformance metric for write I/O operations.

metric

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
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.

Name	Type	Description
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

destination_aggregate

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

Name	Type	Description
cutover_window	integer	Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.
destination_aggregate	destination_aggregate	Aggregate
percent_complete	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.

Name	Type	Description
tiering_policy	string	Tiering policy for FabricPool

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

nas

Name	Type	Description
export_policy	export_policy	Export Policy
gid	integer	The UNIX group ID of the volume. Valid in POST or PATCH.
path	string	The fully-qualified path in the owning SVM’s namespace at which the volume is mounted. The path is case insensitive and must be unique within a SVM’s namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. This attribute is reported in GET only when the volume is mounted.
security_style	string	Security style associated with the volume. Valid in POST or PATCH. mixed ‐ Mixed-style security ntfs ‐ NTFS/WIndows-style security unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions unix ‐ Unix-style security.

Name	Type	Description
uid	integer	The UNIX user ID of the volume. Valid in POST or PATCH.
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 dight), 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	_links	
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. Default is 15000 on AFF platforms and 10000 on all other platforms.

Name	Type	Description
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.
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	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.

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.</num></num></num></num></num>

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

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

snaplock

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

Name	Type	Description
autocommit_period	string	Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num>
compliance_clock_time	string	This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
expiry_time	string	Expiry time of the volume.
is_audit_log	boolean	Indicates if this volume has been configured as SnapLock audit log volume for the SVM .
litigation_count	integer	Litigation count indicates the number of active legal-holds on the volume.

Name	Type	Description
privileged_delete	string	Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled.
retention	retention	
type	string	The SnapLock type of the volume. compliance ‐ A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files. enterprise ‐ An administrator can delete a SnapLock Enterprise(SLE) volume. non_snaplock ‐ Indicates the volume is non-snaplock.

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

logical_space

Name	Type	Description
available	integer	The amount of space available in this volume with storage efficiency space considered used, in bytes.

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	logical_space	
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	snapshot	
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	Peformance 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	Peformance 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	Peformance 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.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

tiering

Name	Type	Description
policy	string	Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH. all ‐ This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks. auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store none ‐ Volume blocks will not be tiered to the cloud store. snapshot_only ‐ This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup.

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	_links	
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
application	application	
autosize	autosize	
clone	clone	
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	efficiency	
encryption	encryption	
error_state	error_state	
files	files	
flexcache_endpoint_type	string	FlexCache endpoint type. none ‐ The volume is neither a FlexCache nor origin of any FlexCache. cache ‐ The volume is a FlexCache volume. origin ‐ The volume is origin of a FlexCache volume.
guarantee	guarantee	
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	metric	Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
movement	movement	Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
name	string	Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.
nas	nas	
qos	qos	QoS information
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.
size	integer	Physical size of the volume. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size.
snaplock	snaplock	
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.
space	space	

Name	Type	Description
state	string	Volume state. A volume can only be brought online if it is offline. 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 ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroups.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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
_links	_links	
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
application	application	

Name	Type	Description
autosize	autosize	
clone	clone	
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	efficiency	
encryption	encryption	
error_state	error_state	
files	files	
flexcache_endpoint_type	string	FlexCache endpoint type. none ‐ The volume is neither a FlexCache nor origin of any FlexCache. cache ‐ The volume is a FlexCache volume. origin ‐ The volume is origin of a FlexCache volume.
guarantee	guarantee	
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	metric	Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
movement	movement	Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
name	string	Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.
nas	nas	
qos	qos	QoS information
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.
size	integer	Physical size of the volume. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size.
snaplock	snaplock	
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.
space	space	

Name	Type	Description
state	string	Volume state. A volume can only be brought online if it is offline. 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 ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroups.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1

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": "string"
  },
  "clone": {
    "parent_snapshot": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "parent_volume": {
```

```
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "split_complete_percent": 0,
    "split_estimate": 0
  },
  "comment": "string",
  "create_time": "2018-06-04 19:00:00 UTC",
  "efficiency": {
    "compaction": "string",
    "compression": "string",
    "cross_volume_dedupe": "string",
    "dedupe": "string"
  },
  "encryption": {
    "key_id": "string",
    "state": "string",
    "status": {
      "code": "string",
      "message": "string"
    },
    "type": "string"
  },
  "files": {
    "used": 0
  },
  "flexcache_endpoint_type": "string",
  "guarantee": {
    "type": "string"
  },
  "language": "string",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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": "string"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 100,
      "name": "default"
    },
    "path": "/user/my_volume",
    "security_style": "string",
    "unix_permissions": 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": "string"
  },
  "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": "string",
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  },
  "style": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "policy": "string"
  },
  "type": "string",
  "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
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	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
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 threhold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size..
maximum	integer	Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume.
minimum	integer	Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume.
mode	string	Autosize mode for the volume. grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value. grow_shrink ‐ Volume grows or shrinks in response to the amount of space used. off ‐ Autosizing of the volume is disabled.

Name	Type	Description
shrink_threshold	integer	Used space threshold size, in percentage, for the automatic shrinkage of the volume. When the amount of used space in the volume drops below this threshold, the volume automatically shrinks unless it has reached the minimum size. The volume shrinks when the 'space.used' is less than the 'shrink_threshold' percent of 'space.size'. The 'shrink_threshold' size cannot be greater than or equal to the 'grow_threshold' size.

snapshot_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

parent_svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

parent_volume

Name	Type	Description
_links	_links	
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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

clone

Name	Type	Description
is_flexclone	boolean	Specifies if this volume is a normal FlexVol or FlexClone. This field needs to be set when creating a FlexClone. Valid in POST.
parent_snapshot	snapshot_reference	
parent_svm	parent_svm	SVM, applies only to SVM-scoped objects.
parent_volume	parent_volume	
split_complete_percent	integer	Percentage of FlexClone blocks split from its parent volume.
split_estimate	integer	Space required by the containing-aggregate to split the FlexClone volume.
split_initiated	boolean	This field is set when split is executed on any FlexClone, that is when the FlexClone volume is split from its parent FlexVol. This field needs to be set for splitting a FlexClone form FlexVol. Valid in PATCH.

efficiency

Name	Type	Description
compaction	string	The system can be enabled/disabled compaction. inline ‐ Data will be compacted first and written to the volume. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.
compression	string	The system can be enabled/disabled compression. inline ‐ Data will be compressed first and written to the volume. background ‐ Data will be written to the volume and compressed later. both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.
cross_volume_dedupe	string	The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. inline ‐ Data will be cross volume deduped first and written to the volume. background ‐ Data will be written to the volume and cross volume deduped later. both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.

Name	Type	Description
dedupe	string	The system can be enabled/disabled dedupe. inline ‐ Data will be deduped first and written to the volume. background ‐ Data will be written to the volume and deduped later. both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.

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.

Name	Type	Description
state	string	Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one.
status	status	
type	string	Volume encryption type. none ‐ The volume is a plain-text one. volume ‐ The volume is encrypted with volume key (NVE volume). aggregate ‐ 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 ‐ File system is inconsistent. false ‐ 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?
type	string	The type of space guarantee of this volume in the aggregate.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Peformance 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	Peformance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Peformance metric for write I/O operations.

metric

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
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.

Name	Type	Description
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

destination_aggregate

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

Name	Type	Description
cutover_window	integer	Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.
destination_aggregate	destination_aggregate	Aggregate
percent_complete	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.

Name	Type	Description
tiering_policy	string	Tiering policy for FabricPool

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

nas

Name	Type	Description
export_policy	export_policy	Export Policy
gid	integer	The UNIX group ID of the volume. Valid in POST or PATCH.
path	string	The fully-qualified path in the owning SVM’s namespace at which the volume is mounted. The path is case insensitive and must be unique within a SVM’s namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. This attribute is reported in GET only when the volume is mounted.
security_style	string	Security style associated with the volume. Valid in POST or PATCH. mixed ‐ Mixed-style security ntfs ‐ NTFS/WIndows-style security unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions unix ‐ Unix-style security.

Name	Type	Description
uid	integer	The UNIX user ID of the volume. Valid in POST or PATCH.
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 dight), 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	_links	
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. Default is 15000 on AFF platforms and 10000 on all other platforms.

Name	Type	Description
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.
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	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.

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.</num></num></num></num></num>

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

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

snaplock

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

Name	Type	Description
autocommit_period	string	Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num>
compliance_clock_time	string	This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
expiry_time	string	Expiry time of the volume.
is_audit_log	boolean	Indicates if this volume has been configured as SnapLock audit log volume for the SVM .
litigation_count	integer	Litigation count indicates the number of active legal-holds on the volume.

Name	Type	Description
privileged_delete	string	Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled.
retention	retention	
type	string	The SnapLock type of the volume. compliance ‐ A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files. enterprise ‐ An administrator can delete a SnapLock Enterprise(SLE) volume. non_snaplock ‐ Indicates the volume is non-snaplock.

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

logical_space

Name	Type	Description
available	integer	The amount of space available in this volume with storage efficiency space considered used, in bytes.

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	logical_space	
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	snapshot	
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	Peformance 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	Peformance 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	Peformance 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.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

tiering

Name	Type	Description
policy	string	Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH. all ‐ This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks. auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store none ‐ Volume blocks will not be tiered to the cloud store. snapshot_only ‐ This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup.

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	_links	
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
application	application	
autosize	autosize	
clone	clone	
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	efficiency	
encryption	encryption	
error_state	error_state	
files	files	
flexcache_endpoint_type	string	FlexCache endpoint type. none ‐ The volume is neither a FlexCache nor origin of any FlexCache. cache ‐ The volume is a FlexCache volume. origin ‐ The volume is origin of a FlexCache volume.
guarantee	guarantee	
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	metric	Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
movement	movement	Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
name	string	Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.
nas	nas	
qos	qos	QoS information
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.
size	integer	Physical size of the volume. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size.
snaplock	snaplock	
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.
space	space	

Name	Type	Description
state	string	Volume state. A volume can only be brought online if it is offline. 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 ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroups.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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	job_link	

Example response


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


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a volume

GET `/storage/volumes/{uuid}`

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
uuid	string	path	True	Unique identifier of the volume.
fields	array[string]	query	False	Specify the fields to return.

Response

```
Status: 200, Ok
```

Name	Type	Description
_links	_links	
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
application	application	
autosize	autosize	
clone	clone	
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	efficiency	
encryption	encryption	
error_state	error_state	
files	files	
flexcache_endpoint_type	string	FlexCache endpoint type. none ‐ The volume is neither a FlexCache nor origin of any FlexCache. cache ‐ The volume is a FlexCache volume. origin ‐ The volume is origin of a FlexCache volume.
guarantee	guarantee	

Name	Type	Description
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	metric	Performance numbers, such as IOPS latency and throughput.
movement	movement	Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
name	string	Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.
nas	nas	
qos	qos	QoS information
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.
size	integer	Physical size of the volume. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size.
snaplock	snaplock	
snapshot_policy	snapshot_policy	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 ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroups.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1

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": "string"
  },
  "clone": {
    "parent_snapshot": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "parent_volume": {
```

```
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "split_complete_percent": 0,
    "split_estimate": 0
  },
  "comment": "string",
  "create_time": "2018-06-04 19:00:00 UTC",
  "efficiency": {
    "compaction": "string",
    "compression": "string",
    "cross_volume_dedupe": "string",
    "dedupe": "string"
  },
  "encryption": {
    "key_id": "string",
    "state": "string",
    "status": {
      "code": "string",
      "message": "string"
    },
    "type": "string"
  },
  "files": {
    "used": 0
  },
  "flexcache_endpoint_type": "string",
  "guarantee": {
    "type": "string"
  },
  "language": "string",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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": "string"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 100,
      "name": "default"
    },
    "path": "/user/my_volume",
    "security_style": "string",
    "unix_permissions": 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": "string"
  },
  "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": "string",
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  },
  "style": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "policy": "string"
  },
  "type": "string",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
```

Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

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 threhold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size..
maximum	integer	Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume.
minimum	integer	Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume.
mode	string	Autosize mode for the volume. grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value. grow_shrink ‐ Volume grows or shrinks in response to the amount of space used. off ‐ Autosizing of the volume is disabled.

Name	Type	Description
shrink_threshold	integer	Used space threshold size, in percentage, for the automatic shrinkage of the volume. When the amount of used space in the volume drops below this threshold, the volume automatically shrinks unless it has reached the minimum size. The volume shrinks when the 'space.used' is less than the 'shrink_threshold' percent of 'space.size'. The 'shrink_threshold' size cannot be greater than or equal to the 'grow_threshold' size.

snapshot_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

parent_svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

parent_volume

Name	Type	Description
_links	_links	
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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

clone

Name	Type	Description
is_flexclone	boolean	Specifies if this volume is a normal FlexVol or FlexClone. This field needs to be set when creating a FlexClone. Valid in POST.
parent_snapshot	snapshot_reference	
parent_svm	parent_svm	SVM, applies only to SVM-scoped objects.
parent_volume	parent_volume	
split_complete_percent	integer	Percentage of FlexClone blocks split from its parent volume.
split_estimate	integer	Space required by the containing-aggregate to split the FlexClone volume.
split_initiated	boolean	This field is set when split is executed on any FlexClone, that is when the FlexClone volume is split from its parent FlexVol. This field needs to be set for splitting a FlexClone form FlexVol. Valid in PATCH.

efficiency

Name	Type	Description
compaction	string	The system can be enabled/disabled compaction. inline ‐ Data will be compacted first and written to the volume. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.
compression	string	The system can be enabled/disabled compression. inline ‐ Data will be compressed first and written to the volume. background ‐ Data will be written to the volume and compressed later. both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.
cross_volume_dedupe	string	The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. inline ‐ Data will be cross volume deduped first and written to the volume. background ‐ Data will be written to the volume and cross volume deduped later. both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.

Name	Type	Description
dedupe	string	The system can be enabled/disabled dedupe. inline ‐ Data will be deduped first and written to the volume. background ‐ Data will be written to the volume and deduped later. both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.

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.

Name	Type	Description
state	string	Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one.
status	status	
type	string	Volume encryption type. none ‐ The volume is a plain-text one. volume ‐ The volume is encrypted with volume key (NVE volume). aggregate ‐ 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 ‐ File system is inconsistent. false ‐ 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?
type	string	The type of space guarantee of this volume in the aggregate.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Peformance 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	Peformance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Peformance metric for write I/O operations.

metric

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
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.

Name	Type	Description
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

destination_aggregate

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

Name	Type	Description
cutover_window	integer	Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.
destination_aggregate	destination_aggregate	Aggregate
percent_complete	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.

Name	Type	Description
tiering_policy	string	Tiering policy for FabricPool

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

nas

Name	Type	Description
export_policy	export_policy	Export Policy
gid	integer	The UNIX group ID of the volume. Valid in POST or PATCH.
path	string	The fully-qualified path in the owning SVM’s namespace at which the volume is mounted. The path is case insensitive and must be unique within a SVM’s namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. This attribute is reported in GET only when the volume is mounted.
security_style	string	Security style associated with the volume. Valid in POST or PATCH. mixed ‐ Mixed-style security ntfs ‐ NTFS/WIndows-style security unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions unix ‐ Unix-style security.

Name	Type	Description
uid	integer	The UNIX user ID of the volume. Valid in POST or PATCH.
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 dight), 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	_links	
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. Default is 15000 on AFF platforms and 10000 on all other platforms.

Name	Type	Description
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.
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	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.

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.</num></num></num></num></num>

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

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

snaplock

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

Name	Type	Description
autocommit_period	string	Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num>
compliance_clock_time	string	This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
expiry_time	string	Expiry time of the volume.
is_audit_log	boolean	Indicates if this volume has been configured as SnapLock audit log volume for the SVM .
litigation_count	integer	Litigation count indicates the number of active legal-holds on the volume.

Name	Type	Description
privileged_delete	string	Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled.
retention	retention	
type	string	The SnapLock type of the volume. compliance ‐ A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files. enterprise ‐ An administrator can delete a SnapLock Enterprise(SLE) volume. non_snaplock ‐ Indicates the volume is non-snaplock.

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

logical_space

Name	Type	Description
available	integer	The amount of space available in this volume with storage efficiency space considered used, in bytes.

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	logical_space	
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	snapshot	
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	Peformance 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	Peformance 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	Peformance 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.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

tiering

Name	Type	Description
policy	string	Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH. all ‐ This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks. auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store none ‐ Volume blocks will not be tiered to the cloud store. snapshot_only ‐ This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup.

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[error_arguments]	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
uuid	string	path	True	Unique identifier of the volume.
restore_to.snapshot.uuid	string	query	False	UUID of the Snapshot copy to restore volume to the point in time the Snapshot copy was taken.
restore_to.snapshot.name	string	query	False	Name of the Snapshot copy to restore volume to the point in time the Snapshot copy was taken.

Name	Type	In	Required	Description
sizing_method	string	query	False	Represents the method to modify the size of a Flexgroup. The following methods are supported: • use_existing_resources - Increases or decreases the size of the FlexGroup by increasing or decreasing the size of the current FlexGroup resources • add_new_resources - Increases the size of the FlexGroup by adding new resources • Default value: 1 • enum: ["use_existing_resources", "add_new_resources"]
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	_links	
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
application	application	
autosize	autosize	
clone	clone	
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	efficiency	
encryption	encryption	
error_state	error_state	
files	files	
flexcache_endpoint_type	string	FlexCache endpoint type. none ‐ The volume is neither a FlexCache nor origin of any FlexCache. cache ‐ The volume is a FlexCache volume. origin ‐ The volume is origin of a FlexCache volume.
guarantee	guarantee	

Name	Type	Description
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	metric	Performance numbers, such as IOPS latency and throughput.
movement	movement	Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
name	string	Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.
nas	nas	
qos	qos	QoS information
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.
size	integer	Physical size of the volume. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size.
snaplock	snaplock	
snapshot_policy	snapshot_policy	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 ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroups.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1

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": "string"
  },
  "clone": {
    "parent_snapshot": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "parent_volume": {
```

```
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "split_complete_percent": 0,
    "split_estimate": 0
  },
  "comment": "string",
  "create_time": "2018-06-04 19:00:00 UTC",
  "efficiency": {
    "compaction": "string",
    "compression": "string",
    "cross_volume_dedupe": "string",
    "dedupe": "string"
  },
  "encryption": {
    "key_id": "string",
    "state": "string",
    "status": {
      "code": "string",
      "message": "string"
    },
    "type": "string"
  },
  "files": {
    "used": 0
  },
  "flexcache_endpoint_type": "string",
  "guarantee": {
    "type": "string"
  },
  "language": "string",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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": "string"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 100,
      "name": "default"
    },
    "path": "/user/my_volume",
    "security_style": "string",
    "unix_permissions": 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": "string"
  },
  "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": "string",
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  },
  "style": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "policy": "string"
  },
  "type": "string",
  "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
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	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
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 threhold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size..
maximum	integer	Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume.
minimum	integer	Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume.
mode	string	Autosize mode for the volume. grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value. grow_shrink ‐ Volume grows or shrinks in response to the amount of space used. off ‐ Autosizing of the volume is disabled.

Name	Type	Description
shrink_threshold	integer	Used space threshold size, in percentage, for the automatic shrinkage of the volume. When the amount of used space in the volume drops below this threshold, the volume automatically shrinks unless it has reached the minimum size. The volume shrinks when the 'space.used' is less than the 'shrink_threshold' percent of 'space.size'. The 'shrink_threshold' size cannot be greater than or equal to the 'grow_threshold' size.

snapshot_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

parent_svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

parent_volume

Name	Type	Description
_links	_links	
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. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

clone

Name	Type	Description
is_flexclone	boolean	Specifies if this volume is a normal FlexVol or FlexClone. This field needs to be set when creating a FlexClone. Valid in POST.
parent_snapshot	snapshot_reference	
parent_svm	parent_svm	SVM, applies only to SVM-scoped objects.
parent_volume	parent_volume	
split_complete_percent	integer	Percentage of FlexClone blocks split from its parent volume.
split_estimate	integer	Space required by the containing-aggregate to split the FlexClone volume.
split_initiated	boolean	This field is set when split is executed on any FlexClone, that is when the FlexClone volume is split from its parent FlexVol. This field needs to be set for splitting a FlexClone form FlexVol. Valid in PATCH.

efficiency

Name	Type	Description
compaction	string	The system can be enabled/disabled compaction. inline ‐ Data will be compacted first and written to the volume. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.
compression	string	The system can be enabled/disabled compression. inline ‐ Data will be compressed first and written to the volume. background ‐ Data will be written to the volume and compressed later. both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.
cross_volume_dedupe	string	The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. inline ‐ Data will be cross volume deduped first and written to the volume. background ‐ Data will be written to the volume and cross volume deduped later. both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.

Name	Type	Description
dedupe	string	The system can be enabled/disabled dedupe. inline ‐ Data will be deduped first and written to the volume. background ‐ Data will be written to the volume and deduped later. both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.

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.

Name	Type	Description
state	string	Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one.
status	status	
type	string	Volume encryption type. none ‐ The volume is a plain-text one. volume ‐ The volume is encrypted with volume key (NVE volume). aggregate ‐ 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 ‐ File system is inconsistent. false ‐ 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?
type	string	The type of space guarantee of this volume in the aggregate.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Peformance 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	Peformance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Peformance metric for write I/O operations.

metric

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
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.

Name	Type	Description
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

destination_aggregate

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

Name	Type	Description
cutover_window	integer	Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.
destination_aggregate	destination_aggregate	Aggregate
percent_complete	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.

Name	Type	Description
tiering_policy	string	Tiering policy for FabricPool

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

nas

Name	Type	Description
export_policy	export_policy	Export Policy
gid	integer	The UNIX group ID of the volume. Valid in POST or PATCH.
path	string	The fully-qualified path in the owning SVM’s namespace at which the volume is mounted. The path is case insensitive and must be unique within a SVM’s namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. This attribute is reported in GET only when the volume is mounted.
security_style	string	Security style associated with the volume. Valid in POST or PATCH. mixed ‐ Mixed-style security ntfs ‐ NTFS/WIndows-style security unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions unix ‐ Unix-style security.

Name	Type	Description
uid	integer	The UNIX user ID of the volume. Valid in POST or PATCH.
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 dight), 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	_links	
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. Default is 15000 on AFF platforms and 10000 on all other platforms.

Name	Type	Description
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.
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	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.

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.</num></num></num></num></num>

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

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

snaplock

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

Name	Type	Description
autocommit_period	string	Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num>
compliance_clock_time	string	This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
expiry_time	string	Expiry time of the volume.
is_audit_log	boolean	Indicates if this volume has been configured as SnapLock audit log volume for the SVM .
litigation_count	integer	Litigation count indicates the number of active legal-holds on the volume.

Name	Type	Description
privileged_delete	string	Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled.
retention	retention	
type	string	The SnapLock type of the volume. compliance ‐ A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files. enterprise ‐ An administrator can delete a SnapLock Enterprise(SLE) volume. non_snaplock ‐ Indicates the volume is non-snaplock.

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

logical_space

Name	Type	Description
available	integer	The amount of space available in this volume with storage efficiency space considered used, in bytes.

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	logical_space	
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	snapshot	
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	Peformance 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	Peformance 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	Peformance 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.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

tiering

Name	Type	Description
policy	string	Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH. all ‐ This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks. auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store none ‐ Volume blocks will not be tiered to the cloud store. snapshot_only ‐ This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup.

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	_links	
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
application	application	
autosize	autosize	
clone	clone	
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	efficiency	
encryption	encryption	
error_state	error_state	
files	files	
flexcache_endpoint_type	string	FlexCache endpoint type. none ‐ The volume is neither a FlexCache nor origin of any FlexCache. cache ‐ The volume is a FlexCache volume. origin ‐ The volume is origin of a FlexCache volume.
guarantee	guarantee	
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	metric	Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
movement	movement	Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
name	string	Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.
nas	nas	
qos	qos	QoS information
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.
size	integer	Physical size of the volume. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size.
snaplock	snaplock	
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.
space	space	

Name	Type	Description
state	string	Volume state. A volume can only be brought online if it is offline. 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 ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroups.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve 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	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1
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
The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1	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
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 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	href	
self	href	

_links

Name	Type	Description
self	href	

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Peformance 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	Peformance 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	Peformance metric for write I/O operations.

records

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	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[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage volume Snapshot copies

Storage volumes volume.uuid snapshots endpoint overview

Overview

A Snapshot copy is the view of the filesystem as it exits 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.

– POST /api/storage/volumes/{volume.uuid}/snapshots

– GET /api/storage/volumes/{volume.uuid}/snapshots

– GET /api/storage/volumes/{volume.uuid}/snapshots/{uuid}

– PATCH /api/storage/volumes/{volume.uuid}/snapshots/{uuid}

– DELETE /api/storage/volumes/{volume.uuid}/snapshots/{uuid}

Examples

Creating a Snapshot copy

The POST operation is used to create a Snapshot copy with the specified attributes.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots

# The call:
curl -X POST "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots" -H "accept:
application/hal+json" -d '{"name": "snapshot_copy", "comment": "Store this
copy." }'

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 13 Mar 2019 22:43:34 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/?name=snapshot_copy
Content-Length: 189
Content-Type: application/json
{
"num_records": 1,
"records": [
  {
    "volume": {
      "name": "v2"
    },
    "svm": {
      "uuid": "8139f958-3c6e-11e9-a45f-005056bbc848",
      "name": "vs0"
    }
    "name": "snapshot_copy",
    "comment": "Store this copy."
  }
],
"job": {
  "uuid": "6f68c85b-45e1-11e9-8fc7-005056bbc848",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/6f68c85b-45e1-11e9-8fc7-005056bbc848"
    }
  }
}
```

```
}

# The Job:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 22:43:57 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 224
Content-Type: application/json
{
"uuid": "6f68c85b-45e1-11e9-8fc7-005056bbc848",
"description": "POST /api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/?name=snapshot_copy",
"state": "success",
"message": "success",
"code": 0
}
```

Retrieving Snapshot copy attributes

The GET operation is used to retrieve Snapshot copy attributes.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots" -H "accept:
application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 21:14:06 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Type: application/json
Transfer-Encoding: chunked
{
"records": [
  {
    "uuid": "402b6c73-73a0-4e89-a58a-75ee0ab3e8c0",
    "name": "hourly.2019-03-13_1305",
    "_links": {
      "self": {
```

```
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0"
      }
    }
  },
  {
    "uuid": "f0dd497f-efe8-44b7-a4f4-bdd3890bc0c8",
    "name": "hourly.2019-03-13_1405",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/f0dd497f-efe8-44b7-a4f4-bdd3890bc0c8"
      }
    }
  },
  {
    "uuid": "02701900-51bd-46b8-9c77-47d9a9e2ce1d",
    "name": "hourly.2019-03-13_1522",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/02701900-51bd-46b8-9c77-47d9a9e2ce1d"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots"
  }
}
}
```

Retrieving the attributes of a specific Snapshot copy

The GET operation is used to retrieve the attributes of a specific Snapshot copy.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-
acb6-005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0" -H
"accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 22:39:26 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 308
Content-Type: application/json
{
"volume": {
  "uuid": "0353dc05-405f-11e9-acb6-005056bbc848",
  "name": "v2",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848"
    }
  }
},
"uuid": "402b6c73-73a0-4e89-a58a-75ee0ab3e8c0",
"svm": {
  "uuid": "8139f958-3c6e-11e9-a45f-005056bbc848",
  "name": "vs0",
  "_links": {
    "self": {
      "href": "/api/svm/svms/8139f958-3c6e-11e9-a45f-005056bbc848"
    }
  }
},
"name": "hourly.2019-03-13_1305",
"create_time": "2019-03-13T13:05:00-04:00",
"_links": {
  "self": {
    "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0"
  }
}
}
```

Updating a Snapshot copy

The PATCH operation is used to update the specific attributes of a Snapshot copy.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-
acb6-005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f" -d
'{"name": "snapshot_copy_new" }' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 13 Mar 2019 22:50:44 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "6f7c3a82-45e2-11e9-8fc7-005056bbc848",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/6f7c3a82-45e2-11e9-8fc7-005056bbc848"
    }
  }
}
}

# The Job:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 22:54:16 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 242
Content-Type: application/json
{
"uuid": "6f7c3a82-45e2-11e9-8fc7-005056bbc848",
"description": "PATCH /api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f",
"state": "success",
"message": "success",
"code": 0
}
```

Deleting a Snapshot copy

The DELETE operation is used to delete a Snapshot copy.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-
acb6-005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f" -H
"accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 13 Mar 2019 22:57:51 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
"job": {
  "uuid": "6da1dfdd-45e3-11e9-8fc7-005056bbc848",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/6da1dfdd-45e3-11e9-8fc7-005056bbc848"
    }
  }
}
}

# The Job:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 23:02:46 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 243
Content-Type: application/json
{
"uuid": "6da1dfdd-45e3-11e9-8fc7-005056bbc848",
"description": "DELETE /api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f",
"state": "success",
"message": "success",
"code": 0
}
```

Retrieve volume Snapshot copies

GET `/storage/volumes/{volume.uuid}/snapshots`

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	_links	
num_records	integer	Number of records
records	array[snapshot]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "create_time": "2019-02-04 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": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

snapshot

The Snapshot copy object represents a point in time Snapshot copy of a volume.

Name	Type	Description
_links	_links	
comment	string	A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.
create_time	string	Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.
expiry_time	string	The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.
name	string	Snapshot copy. Valid in POST or PATCH.
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
_links	_links	
comment	string	A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.
create_time	string	Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.
expiry_time	string	The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.
name	string	Snapshot copy. Valid in POST or PATCH.
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": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```


Response

```
Status: 202, Accepted
```

Name	Type	Description
job	job_link	

Example response


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


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

snapshot

The Snapshot copy object represents a point in time Snapshot copy of a volume.

Name	Type	Description
_links	_links	

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

Delete a volume Snapshot copy

DELETE `/storage/volumes/{volume.uuid}/snapshots/{uuid}`

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	job_link	

Example response

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

Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve volume Snapshot copy details

GET `/storage/volumes/{volume.uuid}/snapshots/{uuid}`

Retrieves details of a specific volume Snapshot copy.

Related ONTAP commands

- `snapshot show`

Learn more

- DOC /storage/volumes/{volume.uuid}/snapshots

Parameters

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

Response

```
Status: 200, Ok
```

Name	Type	Description
_links	_links	
comment	string	A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.
create_time	string	Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.
expiry_time	string	The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.
name	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": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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
_links	_links	
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": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```


Response

```
Status: 202, Accepted
```

Name	Type	Description
job	job_link	

Example response


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


Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error


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


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7

snapshot

The Snapshot copy object represents a point in time Snapshot copy of a volume.

Name	Type	Description
_links	_links	

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

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.