



Manage storage volumes

ONTAP 9.6 REST API reference

NetApp
April 02, 2024

Table of Contents

- Manage storage volumes 1
 - Storage volumes endpoint overview 1
 - Retrieve volumes 16
 - Create a volume on an SVM and storage aggregates 65
 - Delete a volume 107
 - Retrieve a volume 109
 - Update volume attributes 148

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" : "voll", "state" : "online", "type" :
"RW", "aggregates" : [{"name" : "aggr1"}, {"name" : "aggr2"},
{"name":"aggr3"}], "constituents_per_aggregate" : "1", "svm" : {"name" :
"vs1"}, "size" : "240MB", "encryption" : {"enabled" : "False"},
"efficiency" : {"compression" : "both"}, "autosize" : {"maximum" :
"500MB", "minimum" : "240MB"}}'
```

```
# The response:
{
  "job": {
    "uuid": "3cfa38bd-3a78-11e9-ae39-0050568ed7dd",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/3cfa38bd-3a78-11e9-ae39-0050568ed7dd"
      }
    }
  }
}
```

Creating a FlexClone and specifying its properties using POST

```

# The API:
/api/storage/volumes

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name":"voll_clone",{"clone":"parent_volume":
{"name": "voll1"}},{"svm":{"name": "vs0"}, {"clone":
"is_flexclone":"true"}}'

# The response:
HTTP/1.1 202 Accepted
Date: Tue, 26 Feb 2019 09:06:22 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/volumes/?name=voll_clone
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "c9ee0040-39a5-11e9-9b24-00a098439a83",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/c9ee0040-39a5-11e9-9b24-00a098439a83"
      }
    }
  }
}

```

Volumes reported in the GET REST API

The following types of volumes are reported:

- – 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 |
|---|---------|-------|----------|---|
| <code>nas.unix_permissions</code> | integer | query | False | Filter by <code>nas.unix_permissions</code> |
| <code>nas.export_policy.name</code> | string | query | False | Filter by <code>nas.export_policy.name</code> |
| <code>nas.export_policy.id</code> | integer | query | False | Filter by <code>nas.export_policy.id</code> |
| <code>nas.security_style</code> | string | query | False | Filter by <code>nas.security_style</code> |
| <code>nas.gid</code> | integer | query | False | Filter by <code>nas.gid</code> |
| <code>nas.path</code> | string | query | False | Filter by <code>nas.path</code> |
| <code>nas.uid</code> | integer | query | False | Filter by <code>nas.uid</code> |
| <code>error_state.has_bad_blocks</code> | boolean | query | False | Filter by <code>error_state.has_bad_blocks</code> |

| 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_read | integer | query | False | Filter by statistics.latency_read |
| statistics.latency_write | integer | query | False | Filter by statistics.latency_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": "grow"
    },
    "clone": {
      "parent_snapshot": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "this_snapshot",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "parent_svm": {
        "_links": {
          "self": {
```

```

        "href": "/api/resourcelink"
    },
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"parent_volume": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
},
"split_complete_percent": 0,
"split_estimate": 0
},
"comment": "string",
"create_time": "2018-06-04 19:00:00 UTC",
"efficiency": {
    "compaction": "inline",
    "compression": "inline",
    "cross_volume_dedupe": "inline",
    "dedupe": "inline"
},
"encryption": {
    "key_id": "string",
    "state": "encrypted",
    "status": {
        "code": "string",
        "message": "string"
    },
    },
    "type": "none"
},
"files": {
    "used": 0
},
"flexcache_endpoint_type": "none",
"guarantee": {
    "type": "volume"
},
"language": "ar",
"metric": {
    "_links": {
        "self": {

```



```

    "href": "/api/resourcelink"
  }
},
"duration": "PT15S",
"iops": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"latency": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"status": "ok",
"throughput": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 UTC"
},
"movement": {
  "cutover_window": 30,
  "destination_aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "percent_complete": "string",
  "state": "replicating",
  "tiering_policy": "all"
},
"name": "vol_cs_dept",
"nas": {
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,

```

```

    "name": "default"
  },
  "path": "/user/my_volume",
  "security_style": "mixed",
  "unix_permissions": 493
},
"qos": {
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput_iops": 2000,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"quota": {
  "state": "corrupt"
},
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04 19:00:00 UTC",
  "expiry_time": "Wed Sep 5 11:02:42 GMT 2018",
  "is_audit_log": 1,
  "litigation_count": 10,
  "privileged_delete": "enabled",
  "retention": {
    "default": "P30Y",
    "maximum": "P30Y",
    "minimum": "P30Y"
  },
  "type": "enterprise"
},
"snapshot_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}

```

```

},
"space": {
  "available": 0,
  "block_storage_inactive_user_data": 0,
  "capacity_tier_footprint": 0,
  "footprint": 0,
  "logical_space": {
    "available": 0,
    "used_by_afs": 0
  },
  "metadata": 0,
  "over_provisioned": 0,
  "snapshot": {
    "used": 0
  },
  "used": 0
},
"state": "error",
"statistics": {
  "iops_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 UTC"
},
"style": "flexvol",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}

```

```
    },
    "tiering": {
      "policy": "all"
    },
    "type": "rw",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default, Error

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

Example error

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

Definitions

See Definitions

href

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

_links

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

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

snapshot_reference

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

parent_svm

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 | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> example: 028baa66-41bd-11e9-81d5-00a0986138f7 |

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 from 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 | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

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

throughput

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

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

| Name | Type | Description |
|-------|---------|--|
| write | integer | Performance metric for write I/O operations. |

metric

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 digit), first digit is assumed to be zero. |

policy

When "min_throughput_iops", "max_throughput_iops" or "max_throughput_mbps" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

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

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

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

snaplock

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

| Name | Type | Description |
|-----------------------|---------|---|
| autocommit_period | string | <p>Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<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".</p> |
| compliance_clock_time | string | <p>This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.</p> |
| expiry_time | string | <p>Expiry time of the volume.</p> |
| is_audit_log | boolean | <p>Indicates if this volume has been configured as SnapLock audit log volume for the SVM .</p> |
| litigation_count | integer | <p>Litigation count indicates the number of active legal-holds on the volume.</p> |

| Name | Type | Description |
|-------------------|---------------------------|---|
| privileged_delete | string | Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled. |
| retention | 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 | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

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

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

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

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|----------------|----------------|---|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes does not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

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 | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 |

error_arguments

| Name | Type | Description |
|---------|--------|------------------|
| code | string | 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 |
|-------------------------|--|---|
| <code>_links</code> | <code>_links</code> | |
| <code>aggregates</code> | <code>array[aggregates]</code> | Aggregate hosting the volume. Required on POST. |

| Name | Type | Description |
|----------------------------|-----------------------------|--|
| 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 | |
| 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 | <p>Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.</p> |
| uuid | string | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • 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": "grow"
  },
  "clone": {
    "parent_snapshot": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "parent_volume": {
      "_links": {
        "self": {
```



```

        "href": "/api/resourcelink"
    }
},
"name": "volume1",
"uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
},
"split_complete_percent": 0,
"split_estimate": 0
},
"comment": "string",
"create_time": "2018-06-04 19:00:00 UTC",
"efficiency": {
    "compaction": "inline",
    "compression": "inline",
    "cross_volume_dedupe": "inline",
    "dedupe": "inline"
},
"encryption": {
    "key_id": "string",
    "state": "encrypted",
    "status": {
        "code": "string",
        "message": "string"
    },
    "type": "none"
},
"files": {
    "used": 0
},
"flexcache_endpoint_type": "none",
"guarantee": {
    "type": "volume"
},
"language": "ar",
"metric": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "duration": "PT15S",
    "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
    }
},

```

```

    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  },
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": "string",
    "state": "replicating",
    "tiering_policy": "all"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 100,
      "name": "default"
    },
    "path": "/user/my_volume",
    "security_style": "mixed",
    "unix_permissions": 493
  },
  "qos": {
    "policy": {
      "_links": {

```

```

    "self": {
      "href": "/api/resourcelink"
    }
  },
  "max_throughput_iops": 10000,
  "max_throughput_mbps": 500,
  "min_throughput_iops": 2000,
  "name": "performance",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
},
"quota": {
  "state": "corrupt"
},
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04 19:00:00 UTC",
  "expiry_time": "Wed Sep  5 11:02:42 GMT 2018",
  "is_audit_log": 1,
  "litigation_count": 10,
  "privileged_delete": "enabled",
  "retention": {
    "default": "P30Y",
    "maximum": "P30Y",
    "minimum": "P30Y"
  },
  "type": "enterprise"
},
"snapshot_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"space": {
  "available": 0,
  "block_storage_inactive_user_data": 0,
  "capacity_tier_footprint": 0,
  "footprint": 0,
  "logical_space": {
    "available": 0,
    "used_by_afs": 0
  }
}
}

```

```

    },
    "metadata": 0,
    "over_provisioned": 0,
    "snapshot": {
      "used": 0
    },
    "used": 0
  },
  "state": "error",
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  },
  "style": "flexvol",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "policy": "all"
  },
  "type": "rw",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}

```

Response

Status: 202, Accepted

| Name | Type | Description |
|------|----------|-------------|
| job | 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. |

| Error Code | Description |
|------------|---|
| 918243 | The specified Snapshot copy UUID is not correct for the specified Snapshot copy name. |
| 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 threshold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size.. |
| maximum | integer | Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume. |
| minimum | integer | Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume. |
| mode | string | Autosize mode for the volume. grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value. grow_shrink ‐ Volume grows or shrinks in response to the amount of space used. off ‐ Autosizing of the volume is disabled. |

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

snapshot_reference

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

parent_svm

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 | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> example: 028baa66-41bd-11e9-81d5-00a0986138f7 |

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 from 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 | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

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

throughput

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

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

| Name | Type | Description |
|-------|---------|--|
| write | integer | Performance metric for write I/O operations. |

metric

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 digit), first digit is assumed to be zero. |

policy

When "min_throughput_iops", "max_throughput_iops" or "max_throughput_mbps" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

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

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

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

snaplock

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

| Name | Type | Description |
|-----------------------|---------|---|
| autocommit_period | string | <p>Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<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".</p> |
| compliance_clock_time | string | <p>This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.</p> |
| expiry_time | string | <p>Expiry time of the volume.</p> |
| is_audit_log | boolean | <p>Indicates if this volume has been configured as SnapLock audit log volume for the SVM .</p> |
| litigation_count | integer | <p>Litigation count indicates the number of active legal-holds on the volume.</p> |

| Name | Type | Description |
|-------------------|---------------------------|---|
| privileged_delete | string | Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled. |
| retention | 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 | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

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

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

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

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|----------------|--------------------------------|---|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes does not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

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 | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 |

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 |
|---------------------|----------------------------|--------------------|----------|----------------------------------|
| <code>uuid</code> | <code>string</code> | <code>path</code> | True | Unique identifier of the volume. |
| <code>fields</code> | <code>array[string]</code> | <code>query</code> | False | Specify the fields to return. |

Response

```
Status: 200, Ok
```

| Name | Type | Description |
|----------------------------|-------------------------------------|--|
| _links | _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 | <p>Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.</p> |
| uuid | string | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • 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": "grow"
  },
  "clone": {
    "parent_snapshot": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "parent_volume": {
      "_links": {
        "self": {
```

```

        "href": "/api/resourcelink"
    }
},
"name": "volume1",
"uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
},
"split_complete_percent": 0,
"split_estimate": 0
},
"comment": "string",
"create_time": "2018-06-04 19:00:00 UTC",
"efficiency": {
    "compaction": "inline",
    "compression": "inline",
    "cross_volume_dedupe": "inline",
    "dedupe": "inline"
},
"encryption": {
    "key_id": "string",
    "state": "encrypted",
    "status": {
        "code": "string",
        "message": "string"
    },
    "type": "none"
},
"files": {
    "used": 0
},
"flexcache_endpoint_type": "none",
"guarantee": {
    "type": "volume"
},
"language": "ar",
"metric": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "duration": "PT15S",
    "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
    }
},

```

```

    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  },
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": "string",
    "state": "replicating",
    "tiering_policy": "all"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 100,
      "name": "default"
    },
    "path": "/user/my_volume",
    "security_style": "mixed",
    "unix_permissions": 493
  },
  "qos": {
    "policy": {
      "_links": {

```



```

    "self": {
      "href": "/api/resourcelink"
    }
  },
  "max_throughput_iops": 10000,
  "max_throughput_mbps": 500,
  "min_throughput_iops": 2000,
  "name": "performance",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
},
"quota": {
  "state": "corrupt"
},
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04 19:00:00 UTC",
  "expiry_time": "Wed Sep 5 11:02:42 GMT 2018",
  "is_audit_log": 1,
  "litigation_count": 10,
  "privileged_delete": "enabled",
  "retention": {
    "default": "P30Y",
    "maximum": "P30Y",
    "minimum": "P30Y"
  },
  "type": "enterprise"
},
"snapshot_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"space": {
  "available": 0,
  "block_storage_inactive_user_data": 0,
  "capacity_tier_footprint": 0,
  "footprint": 0,
  "logical_space": {
    "available": 0,
    "used_by_afs": 0
  }
}
}

```

```

    },
    "metadata": 0,
    "over_provisioned": 0,
    "snapshot": {
      "used": 0
    },
    "used": 0
  },
  "state": "error",
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  },
  "style": "flexvol",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "policy": "all"
  },
  "type": "rw",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}

```

Error

Status: Default, Error

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

Example error

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

Definitions

See Definitions

href

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

_links

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

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

snapshot_reference

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

parent_svm

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 | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> example: 028baa66-41bd-11e9-81d5-00a0986138f7 |

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 from 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 | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

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

throughput

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

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

| Name | Type | Description |
|-------|---------|--|
| write | integer | Performance metric for write I/O operations. |

metric

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 digit), first digit is assumed to be zero. |

policy

When "min_throughput_iops", "max_throughput_iops" or "max_throughput_mbps" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

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

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

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

snaplock

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

| Name | Type | Description |
|-----------------------|---------|---|
| autocommit_period | string | <p>Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<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".</p> |
| compliance_clock_time | string | <p>This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.</p> |
| expiry_time | string | <p>Expiry time of the volume.</p> |
| is_audit_log | boolean | <p>Indicates if this volume has been configured as SnapLock audit log volume for the SVM .</p> |
| litigation_count | integer | <p>Litigation count indicates the number of active legal-holds on the volume.</p> |

| Name | Type | Description |
|-------------------|---------------------------|---|
| privileged_delete | string | Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled. |
| retention | 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 | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

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

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

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

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|----------------|----------------|--|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | <p>Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data".</p> <p>"Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated.</p> <p>"Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes does not have the latest data.</p> |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

svm

SVM containing the volume. Required on POST.

| Name | Type | Description |
|------------------------|------------------------|-----------------------------------|
| _links | _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 | <p>Represents the method to modify the size of a Flexgroup. The following methods are supported:</p> <ul style="list-style-type: none"> • use_existing_resources - Increases or decreases the size of the FlexGroup by increasing or decreasing the size of the current FlexGroup resources • add_new_resources - Increases the size of the FlexGroup by adding new resources • Default value: 1 • enum: ["use_existing_resources", "add_new_resources"] |
| 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 | <p>Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.</p> |
| uuid | string | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • 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": "grow"
  },
  "clone": {
    "parent_snapshot": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "parent_volume": {
      "_links": {
        "self": {
```

```

        "href": "/api/resourcelink"
    }
},
"name": "volume1",
"uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
},
"split_complete_percent": 0,
"split_estimate": 0
},
"comment": "string",
"create_time": "2018-06-04 19:00:00 UTC",
"efficiency": {
    "compaction": "inline",
    "compression": "inline",
    "cross_volume_dedupe": "inline",
    "dedupe": "inline"
},
"encryption": {
    "key_id": "string",
    "state": "encrypted",
    "status": {
        "code": "string",
        "message": "string"
    },
    "type": "none"
},
"files": {
    "used": 0
},
"flexcache_endpoint_type": "none",
"guarantee": {
    "type": "volume"
},
"language": "ar",
"metric": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "duration": "PT15S",
    "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
    }
},

```

```

    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  },
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": "string",
    "state": "replicating",
    "tiering_policy": "all"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 100,
      "name": "default"
    },
    "path": "/user/my_volume",
    "security_style": "mixed",
    "unix_permissions": 493
  },
  "qos": {
    "policy": {
      "_links": {

```

```
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "max_throughput_iops": 10000,
  "max_throughput_mbps": 500,
  "min_throughput_iops": 2000,
  "name": "performance",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
},
"quota": {
  "state": "corrupt"
},
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04 19:00:00 UTC",
  "expiry_time": "Wed Sep 5 11:02:42 GMT 2018",
  "is_audit_log": 1,
  "litigation_count": 10,
  "privileged_delete": "enabled",
  "retention": {
    "default": "P30Y",
    "maximum": "P30Y",
    "minimum": "P30Y"
  },
  "type": "enterprise"
},
"snapshot_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"space": {
  "available": 0,
  "block_storage_inactive_user_data": 0,
  "capacity_tier_footprint": 0,
  "footprint": 0,
  "logical_space": {
    "available": 0,
    "used_by_afs": 0
  }
}
```

```

    },
    "metadata": 0,
    "over_provisioned": 0,
    "snapshot": {
      "used": 0
    },
    "used": 0
  },
  "state": "error",
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  },
  "style": "flexvol",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "policy": "all"
  },
  "type": "rw",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}

```


Response

Status: 202, Accepted

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

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

snapshot_reference

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

parent_svm

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 | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> example: 028baa66-41bd-11e9-81d5-00a0986138f7 |

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 from 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 | Performance metric for write I/O operations. |

latency

The round trip latency in microseconds observed at the storage object.

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

throughput

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

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

| Name | Type | Description |
|-------|---------|--|
| write | integer | Performance metric for write I/O operations. |

metric

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 digit), first digit is assumed to be zero. |

policy

When "min_throughput_iops", "max_throughput_iops" or "max_throughput_mbps" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

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

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

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

snaplock

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

| Name | Type | Description |
|-----------------------|---------|---|
| autocommit_period | string | <p>Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<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".</p> |
| compliance_clock_time | string | <p>This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.</p> |
| expiry_time | string | <p>Expiry time of the volume.</p> |
| is_audit_log | boolean | <p>Indicates if this volume has been configured as SnapLock audit log volume for the SVM .</p> |
| litigation_count | integer | <p>Litigation count indicates the number of active legal-holds on the volume.</p> |

| Name | Type | Description |
|-------------------|---------------------------|---|
| privileged_delete | string | Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled. |
| retention | 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 | Performance metric for write I/O operations. |

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

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

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

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

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

| Name | Type | Description |
|----------------|--------------------------------|---|
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation. |
| status | string | Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes does not have the latest data. |
| throughput_raw | throughput_raw | Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time. |
| timestamp | string | The timestamp of the performance data. |

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 | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 |

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 © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.