



Manage storage volumes

ONTAP 9.10.1 REST API Documentation

NetApp
April 02, 2024

Table of Contents

- Manage storage volumes 1
 - Storage volumes endpoint overview 1
 - Retrieve volumes 22
 - Create a volume on an SVM and storage aggregates 127
 - Delete a volume 205
 - Retrieve a volume 208
 - Update volume attributes 278

Manage storage volumes

Storage volumes endpoint overview

Overview

FlexVol volumes are logical containers used by ONTAP to serve data to clients. They contain file systems in a NAS environment and LUNs in a SAN environment.

A FlexGroup volume is a scale-out NAS container that provides high performance along with automatic load distribution and scalability. A FlexGroup volume contains several constituents that automatically and transparently share the traffic.

FlexClone volumes are writable, point-in-time copies of a FlexVol volume. At this time, FlexClones of FlexGroups are not supported.

Volumes with SnapLock type Compliance or Enterprise, are referred to as SnapLock volumes. Volumes with SnapLock type cannot be of FlexGroup style. Once a SnapLock aggregate is created, by default, volumes created inside the aggregate inherit the "snaplock" property from the aggregate. It is possible to create a SnapLock volume by specifying SnapLock parameters. SnapLock parameters are only available at the "advanced" privilege level.

ONTAP storage APIs allow you to create, modify, and monitor volumes and aggregates.

Storage efficiency

Storage efficiency is used to remove duplicate blocks in the data and to compress the data. Efficiency has deduplication, compression, cross volume deduplication, and compaction options. On All Flash systems, all efficiencies are enabled by default on volume creation. Options such as "background/inline/both" are treated as both, which means both background and inline are enabled for any efficiency option. The option "none" disables both background and inline efficiency.

To enable any efficiency option on all-flash or FAS systems, background deduplication is always enabled.

Quotas

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

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

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

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

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

File System Analytics

File system analytics provide a quick method for obtaining information summarizing properties of all files within any directory tree of a volume. For more information on file system analytics, see [DOC](#) [/storage/volumes{volume.uuid}/files/{path}](#) . Analytics can be enabled or disabled on individual volumes.

The following APIs can be used to enable or disable and obtain analytics state for a FlexVol volume or a FlexGroup volume:

– PATCH /api/storage/volumes/{uuid} -d '{"analytics.state":"on"}

– PATCH /api/storage/volumes/{uuid} -d '{"analytics.state":"off"}

– GET /api/storage/volumes/{uuid}?fields=analytics

QoS

QoS policy and settings enforce Service Level Objectives (SLO) on a volume. SLO can be set by specifying `qos.max_throughput_iops` and/or `qos.max_throughput_mbps` or `qos.min_throughput_iops` and/or `qos.min_throughput_mbps`. Specifying `min_throughput_iops` or `min_throughput_mbps` is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying `qos.name` or `qos.uuid` property.

Performance monitoring

Performance of a volume can be monitored by the `metric.*` and `statistics.*` fields. These show the performance of the volume in terms of IOPS, latency and throughput. The `metric.*` fields denote an average whereas `statistics.*` fields denote a real-time monotonically increasing value aggregated across all nodes.

Volume APIs

The following APIs are used to perform operations related with FlexVol volumes and FlexGroup volumes:

– POST /api/storage/volumes

– GET /api/storage/volumes

– GET /api/storage/volumes/{uuid}

– PATCH /api/storage/volumes/{uuid}

– DELETE /api/storage/volumes/{uuid}

Examples

Creating a volume

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

```

# The API:
/api/storage/volumes

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

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

```

Creating a SnapLock volume and specifying its properties using POST

```

# The API:
/api/storage/volumes

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name": "vol1", "aggregates":[{"name":
"aggr1"}], "svm":{"name" : "vs1"}, "snaplock":{"retention":{"default":
"P20Y"}, "type": "compliance"}}'

# The response:
{
"job": {
  "uuid": "e45b123b-c228-11e8-aa20-0050568e36bb",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/e45b123b-c228-11e8-aa20-0050568e36bb"
    }
  }
}
}

```

Creating a FlexGroup volume and specifying its properties using POST

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

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

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

```

Volumes reported in the GET REST API

The following types of volumes are reported:

- – RW, DP and LS volumes
- – FlexGroup volume
- – FlexCache volume
- – FlexClone volume
- – FlexGroup constituent

The following volumes are not reported:

- – DEL and TMP type volume
- – Node Root volume

– System Vserver volume

– FlexCache constituent

Examples

Retrieving the list of volumes

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

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

# The response:
{
  "records": [
    {
      "uuid": "2d1167cc-c3f2-495a-a23f-8f50b071b9b8",
      "name": "vsdata_root",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/2d1167cc-c3f2-495a-a23f-
8f50b071b9b8"
        }
      }
    },
    {
      "uuid": "3969be7e-78b4-4b4c-82a4-fa86331f03df",
      "name": "vsfg_root",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/3969be7e-78b4-4b4c-82a4-
fa86331f03df"
        }
      }
    },
    {
      "uuid": "59c03ac5-e708-4ce8-a676-278dc249fda2",
      "name": "svm_root",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/59c03ac5-e708-4ce8-a676-
278dc249fda2"
        }
      }
    }
  ]
}
```



```

},
{
  "uuid": "6802635b-8036-11e8-aae5-0050569503ac",
  "name": "fgvol",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/6802635b-8036-11e8-aae5-0050569503ac"
    }
  }
},
{
  "uuid": "d0c3359c-5448-4a9b-a077-e3295a7e9057",
  "name": "datavol",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/d0c3359c-5448-4a9b-a077-e3295a7e9057"
    }
  }
}
],
"num_records": 5,
"_links": {
  "self": {
    "href": "/api/storage/volumes"
  }
}
}
}

```

Retrieving the attributes of a volume

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

```

# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-a077-e3295a7e9057" -H "accept: application/hal+json"

# The response:
{
  "uuid": "d0c3359c-5448-4a9b-a077-e3295a7e9057",
  "comment": "This is a data volume",
  "create_time": "2018-07-05T14:56:44+05:30",

```

```
"language": "en_us",
"name": "datavol",
"size": 20971520,
"state": "online",
"style": "flexvol",
"tiering_policy": "auto",
"type": "rw",
"aggregates": [
  {
    "name": "data",
    "uuid": "aa742322-36bc-4d98-bbc4-0a827534c035",
    "_links": {
      "self": {
        "href": "/api/cluster/aggregates/data"
      }
    }
  }
],
"encryption": {
  "enabled": false,
  "state": "none",
  "key_id": "",
  "type": "none"
},
"error_state": {
  "has_bad_blocks": false,
  "is_inconsistent": false
},
"files": {
  "maximum": 566,
  "used": 96
},
"nas": {
  "gid": 2468,
  "security_style": "unix",
  "uid": 1357,
  "unix_permissions": 4755,
  "export_policy": {
    "name": "default",
    "id": 8589934593
  }
},
"junction_parent": {
  "name": "voll",
  "uuid": "a2564f80-25fb-41e8-9b49-44de2600991f",
  "_links": {
    "self": {
```

```
    "href": "/api/storage/volumes/a2564f80-25fb-41e8-9b49-44de2600991f"
  }
}
},
"metric": {
  "timestamp": "2019-04-09T05:50:15Z",
  "status": "ok",
  "duration": "PT15S",
  "latency": {
    "other": 0,
    "total": 0,
    "read": 0,
    "write": 0
  },
  "iops": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "throughput": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "cloud": {
    "timestamp": "2019-04-09T05:50:15Z",
    "status": "ok",
    "duration": "PT15S",
    "iops" : {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "latency": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    }
  },
  "flexcache": {
```

```
    "timestamp": "2019-04-09T05:50:15Z",
    "status": "ok",
    "duration": "PT1D",
    "cache_miss_percent": 0,
    "bandwidth_savings": 0
  }
},
"statistics": {
  "timestamp": "2019-04-09T05:50:42Z",
  "status": "ok",
  "latency_raw": {
    "other": 38298,
    "total": 38298,
    "read": 0,
    "write": 0
  },
  "iops_raw": {
    "read": 0,
    "write": 0,
    "other": 3,
    "total": 3
  },
  "throughput_raw": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "cloud": {
    "timestamp": "2019-04-09T05:50:42Z",
    "status": "ok",
    "iops_raw" : {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "latency_raw": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    }
  },
  "flexcache_raw": {
    "timestamp": "2019-04-09T05:50:15Z",
```

```
    "status": "ok",
    "cache_miss_blocks": 0,
    "client_requested_blocks": 0
  }
},
"qos": {
  "policy": {
    "min_throughput_iops": 0,
    "min_throughput_mbps": 0,
    "max_throughput_iops": 1000,
    "max_throughput_mbps": 0,
    "uuid": "228454af-5a8b-11e9-bd5b-005056ac6f1f",
    "name": "pg1"
  }
},
"snaplock": {
  "append_mode_enabled": false,
  "autocommit_period": "none",
  "compliance_clock_time": "2019-05-24T10:59:00+05:30",
  "expiry_time": "2038-01-19T08:44:28+05:30",
  "is_audit_log": false,
  "litigation_count": 0,
  "privileged_delete": "disabled",
  "type": "enterprise",
  "retention": {
    "default": "P0Y",
    "minimum": "P0Y",
    "maximum": "P30Y"
  }
},
"snapshot_policy": {
  "name": "default"
},
"svm": {
  "name": "vsdata",
  "uuid": "d61b69f5-7458-11e8-ad3f-0050569503ac"
},
"anti_ransomware_state": "disabled",
"_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/"
    }
  }
}
```

Retrieving the constituents of a FlexGroup volume

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

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-a077-e3295a7e9057/?fields=constituents" -H "accept: application/hal+json"

# The response:
{
  "uuid": "d0c3359c-5448-4a9b-a077-e3295a7e9057",
  "name": "fg",
  "constituents": [
    {
      "name": "fg__0001",
      "space": {
        "size": 536870912,
        "available": 479756288,
        "used": 30273536,
        "local_tier_footprint": 540958720,
        "footprint": 389120,

```

```

"over_provisioned": 0,
"metadata": 4087808,
"total_footprint": 540958720,
"logical_space": {
  "reporting": false,
  "enforcement": false,
  "used_by_afs": 30273536
},
"snapshot": {
  "reserve_percent": 5,
  "autodelete_enabled": false
}
},
"aggregates": {
  "name": "vs_aggr",
  "uuid": "2e00716c-eb54-45a9-8ca5-be50ccd40708",
  "_links": {
    "self": {
      "href": "/api/storage/aggregates/2e00716c-eb54-45a9-8ca5-
be50ccd40708"
    }
  }
}
},
{
  "name": "fg__0002",
  "space": {
    "size": 536870912,
    "available": 479756288,
    "used": 30273536,
    "local_tier_footprint": 540958720,
    "footprint": 389120,
    "over_provisioned": 0,
    "metadata": 4087808,
    "total_footprint": 540958720,
    "logical_space": {
      "reporting": false,
      "enforcement": false,
      "used_by_afs": 30273536
    },
    "snapshot": {
      "reserve_percent": 5,
      "autodelete_enabled": false
    }
  },
  "aggregates": {

```

```

    "name": "aggr1",
    "uuid": "f655cf4f-8208-4fc8-bfc8-4f238d434402",
    "_links": {
      "self": {
        "href": "/api/storage/aggregates/f655cf4f-8208-4fc8-bfc8-4f238d434402"
      }
    }
  },
  {
    "name": "fg__0003",
    "space": {
      "size": 536870912,
      "available": 479756288,
      "used": 30273536,
      "local_tier_footprint": 540958720,
      "footprint": 389120,
      "over_provisioned": 0,
      "metadata": 4087808,
      "total_footprint": 540958720,
      "logical_space": {
        "reporting": false,
        "enforcement": false,
        "used_by_afs": 30273536
      },
      "snapshot": {
        "reserve_percent": 5,
        "autodelete_enabled": false
      }
    },
    "aggregates": {
      "name": "aggr2",
      "uuid": "c7239f14-b861-46fc-b406-70fe13c1a4fb",
      "_links": {
        "self": {
          "href": "/api/storage/aggregates/c7239f14-b861-46fc-b406-70fe13c1a4fb"
        }
      }
    }
  },
  {
    "name": "fg__0004",
    "space": {
      "size": 536870912,

```



```

"available": 479756288,
"used": 30273536,
"local_tier_footprint": 540958720,
"footprint": 389120,
"over_provisioned": 0,
"metadata": 4087808,
"total_footprint": 540958720,
"logical_space": {
  "reporting": false,
  "enforcement": false,
  "used_by_afs": 30273536
},
"snapshot": {
  "reserve_percent": 5,
  "autodelete_enabled": false
},
"aggregates": {
  "name": "aggr3",
  "uuid": "4c86137a-06af-482d-a41c-d64acc5dcea4",
  "_links": {
    "self": {
      "href": "/api/storage/aggregates/4c86137a-06af-482d-a41c-
d64acc5dcea4"
    }
  }
},
],
"_links": {
  "self": {
    "href": "/api/storage/volumes/d0c3359c-5448-4a9b-a077-e3295a7e9057"
  }
}
}

```

Updating the attributes of a volume

Examples

Updating the attributes of a volume

The PATCH request is used to update the attributes of a volume.

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{ "size": 26214400, "nas": {"security_style":
"mixed"}, "comment": "This is a data volume" }' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Tue, 31 Jul 2018 09:36:43 GMT
Server: libzapid-httpd
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "3c5be5a6-94a5-11e8-8ca3-00505695c11b",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/3c5be5a6-94a5-11e8-8ca3-00505695c11b"
      }
    }
  }
}
```

Updating the attributes of a FlexClone using PATCH

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"clone":{"split_initiated":"true"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 25 Feb 2019 10:10:19 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "8e01747f-38e5-11e9-8a3a-00a09843994b",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/8e01747f-38e5-11e9-8a3a-00a09843994b"
      }
    }
  }
}
```

Enabling quotas for a FlexVol or a FlexGroup volume using PATCH

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"quota":{"enabled":"true"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 25 Feb 2019 10:10:19 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "d2fe7299-57d0-11e9-a2dc-005056a7f717",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/d2fe7299-57d0-11e9-a2dc-005056a7f717"
      }
    }
  }
}
```

Disabling quotas for a FlexVol or a FlexGroup volume using PATCH

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"quota":{"enabled":"false"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 25 Feb 2019 10:10:19 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "0c8f6bea-57d1-11e9-a2dc-005056a7f717",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/0c8f6bea-57d1-11e9-a2dc-005056a7f717"
      }
    }
  }
}
```

Add tiering object tags for a FlexVol using PATCH

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"tiering.object_tags": [ "key1=val1", "key2=val2"
]}' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Tue, 11 Feb 2020 19:29:25 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "d05012de-4d04-11ea-836b-005056bb6f9d",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/d05012de-4d04-11ea-836b-005056bb6f9d"
      }
    }
  }
}
```

Remove tiering object tags for a FlexVol using PATCH

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"tiering.object_tags": []}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 Jan 2020 22:28:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "ca234df1-3ef8-11ea-9a56-005056bb69a1",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/ca234df1-3ef8-11ea-9a56-005056bb69a1"
      }
    }
  }
}
```

Deleting a volume

Example

Deleting a volume

The DELETE request is used to delete a volume.

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

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

# The response:
HTTP/1.1 202 Accepted
cache-control: no-cache,no-store,must-revalidate
connection: Keep-Alive
content-length: 189
content-type: application/json
date: Wed, 01 Aug 2018 09:40:36 GMT
keep-alive: timeout=5, max=100
server: libzapid-httpd
{
  "job": {
    "uuid": "f1aa3eb8-956e-11e8-86bf-0050568e2249",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f1aa3eb8-956e-11e8-86bf-0050568e2249"
      }
    }
  }
}
```

Retrieve volumes

GET /storage/volumes

Introduced In: 9.6

Retrieves volumes.

Expensive properties

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

- `is_svm_root`
- `analytics.*`
- `anti_ransomware.*`
- `application.*`

- encryption.*
- queue_for_encryption
- convert_unicode
- 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.*
- max_dir_size
- nas.export_policy.id
- nas.gid
- nas.path
- nas.security_style
- nas.uid
- nas.unix_permissions
- nas.junction_parent.name
- nas.junction_parent.uuid
- snaplock.*
- restore_to.*
- snapshot_policy.uuid
- quota.*
- qos.*
- flexcache_endpoint_type
- space.block_storage_inactive_user_data
- space.capacity_tier_footprint
- space.performance_tier_footprint
- space.local_tier_footprint

- space.footprint
- space.over_provisioned
- space.metadata
- space.total_footprint
- space.dedupe_metafiles_footprint
- space.dedupe_metafiles_temporary_footprint
- space.delayed_free_footprint
- space.file_operation_metadata
- space.snapmirror_destination_footprint
- space.volume_guarantee_footprint
- space.cross_volume_dedupe_metafiles_footprint
- space.cross_volume_dedupe_metafiles_temporary_footprint
- space.snapshot_reserve_unusable
- space.snapshot_spill
- space.user_data
- space.logical_space.*
- space.snapshot.*
- space.used_by_afs
- space.afs_total
- space.available_percent
- space.full_threshold_percent
- space.nearly_full_threshold_percent
- space.overwrite_reserve
- space.overwrite_reserve_used
- space.size_available_for_snapshots
- space.percent_used
- space.fractional_reserve
- space.block_storage_inactive_user_data_percent
- space.physical_used
- space.physical_used_percent
- space.expected_available
- space.filesystem_size
- space.filesystem_size_fixed
- guarantee.*

- autosize.*
- movement.*
- statistics.*
- constituents.name
- constituents.space.size
- constituents.space.available
- constituents.space.used
- constituents.space.available_percent
- constituents.space.used_percent
- constituents.space.block_storage_inactive_user_data
- constituents.space.capacity_tier_footprint
- constituents.space.performance_tier_footprint
- constituents.space.local_tier_footprint
- constituents.space.footprint
- constituents.space.over_provisioned
- constituents.space.metadata
- constituents.space.total_footprint
- constituents.space.logical_space.reporting
- constituents.space.logical_space.enforcement
- constituents.space.logical_space.used_by_afs
- constituents.space.logical_space.available
- constituents.space.snapshot.used
- constituents.space.snapshot.reserve_percent
- constituents.space.snapshot.autodelete_enabled
- constituents.aggregates.name
- constituents.aggregates.uuid
- constituents.movement.destination_aggregate.name
- constituents.movement.destination_aggregate.uuid
- constituents.movement.state
- constituents.movement.percent_complete
- constituents.movement.cutover_window
- constituents.movement.tiering_policy

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
- security anti-ransomware volume show
- security anti-ransomware volume space show

Parameters

| Name | Type | In | Required | Description |
|------------------------|---------|-------|----------|---|
| is_constituent | boolean | query | False | When set to false, only FlexVol and FlexGroup volumes are returned. When set to true, only FlexGroup constituent volumes are returned. Default for GET calls is false. <ul style="list-style-type: none">• Default value:• Introduced in: 9.10 |
| nas.unix_permissions | integer | query | False | Filter by nas.unix_permissions |
| nas.security_style | string | query | False | Filter by nas.security_style |
| nas.export_policy.name | string | query | False | Filter by nas.export_policy.name |

| Name | Type | In | Required | Description |
|-------------------------------------|---------|-------|----------|--|
| nas.export_policy.id | integer | query | False | Filter by nas.export_policy.id |
| nas.junction_parent.uuid | string | query | False | Filter by nas.junction_parent.uuid • Introduced in: 9.9 |
| nas.junction_parent.name | string | query | False | Filter by nas.junction_parent.name • Introduced in: 9.9 |
| nas.uid | integer | query | False | Filter by nas.uid |
| nas.path | string | query | False | Filter by nas.path |
| nas.gid | integer | query | False | Filter by nas.gid |
| movement.start_time | string | query | False | Filter by movement.start_time • Introduced in: 9.9 |
| movement.cutover_window | integer | query | False | Filter by movement.cutover_window |
| movement.percent_complete | integer | query | False | Filter by movement.percent_complete |
| movement.state | string | query | False | Filter by movement.state |
| movement.destination_aggregate.uuid | string | query | False | Filter by movement.destination_aggregate.uuid |

| Name | Type | In | Required | Description |
|---|---------|-------|----------|--|
| movement.destination_aggregate.name | string | query | False | Filter by movement.destination_aggregate.name |
| is_svm_root | boolean | query | False | Filter by is_svm_root <ul style="list-style-type: none"> • Introduced in: 9.7 |
| uuid | string | query | False | Filter by uuid |
| flexcache_endpoint_type | string | query | False | Filter by flexcache_endpoint_type |
| status | string | query | False | Filter by status <ul style="list-style-type: none"> • Introduced in: 9.9 |
| error_state.is_inconsistent | boolean | query | False | Filter by error_state.is_inconsistent |
| error_state.has_bad_blocks | boolean | query | False | Filter by error_state.has_bad_blocks |
| constituents.space.performance_tier_footprint | integer | query | False | Filter by constituents.space.performance_tier_footprint <ul style="list-style-type: none"> • Introduced in: 9.9 |
| constituents.space.available_percent | integer | query | False | Filter by constituents.space.available_percent <ul style="list-style-type: none"> • Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|--|---------|-------|----------|--|
| constituents.space.capacity_tier_footprint | integer | query | False | Filter by constituents.space.capacity_tier_footprint • Introduced in: 9.9 |
| constituents.space.used_percent | integer | query | False | Filter by constituents.space.used_percent • Introduced in: 9.10 |
| constituents.space.used_by_afs | integer | query | False | Filter by constituents.space.used_by_afs • Introduced in: 9.9 |
| constituents.space.available | integer | query | False | Filter by constituents.space.available • Introduced in: 9.9 |
| constituents.space.size | integer | query | False | Filter by constituents.space.size • Introduced in: 9.9 |
| constituents.space.metadata | integer | query | False | Filter by constituents.space.metadata • Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---------|-------|----------|---|
| constituents.space.block_storage_inactive_user_data | integer | query | False | Filter by constituents.space.block_storage_inactive_user_data • Introduced in: 9.9 |
| constituents.space.footprint | integer | query | False | Filter by constituents.space.footprint • Introduced in: 9.9 |
| constituents.space.over_provisioned | integer | query | False | Filter by constituents.space.over_provisioned • Introduced in: 9.9 |
| constituents.space.total_footprint | integer | query | False | Filter by constituents.space.total_footprint • Introduced in: 9.9 |
| constituents.space.local_tier_footprint | integer | query | False | Filter by constituents.space.local_tier_footprint • Introduced in: 9.9 |
| constituents.space.aifs_total | integer | query | False | Filter by constituents.space.aifs_total • Introduced in: 9.9 |
| constituents.space.used | integer | query | False | Filter by constituents.space.used • Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|--|---------|-------|----------|--|
| constituents.space.logical_space.enforcement | boolean | query | False | Filter by constituents.space.logical_space.enforcement • Introduced in: 9.9 |
| constituents.space.logical_space.reporting | boolean | query | False | Filter by constituents.space.logical_space.reporting • Introduced in: 9.9 |
| constituents.space.logical_space.available | integer | query | False | Filter by constituents.space.logical_space.available • Introduced in: 9.9 |
| constituents.space.logical_space.used_by_afs | integer | query | False | Filter by constituents.space.logical_space.used_by_afs • Introduced in: 9.9 |
| constituents.space.snapshot.reserve_percent | integer | query | False | Filter by constituents.space.snapshot.reserve_percent • Introduced in: 9.9 |
| constituents.space.snapshot.used | integer | query | False | Filter by constituents.space.snapshot.used • Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|--|---------|-------|----------|--|
| constituents.aggregates.uuid | string | query | False | Filter by constituents.aggregates.uuid • Introduced in: 9.9 |
| constituents.aggregates.name | string | query | False | Filter by constituents.aggregates.name • Introduced in: 9.9 |
| constituents.movement.cutover_window | integer | query | False | Filter by constituents.movement.cutover_window • Introduced in: 9.9 |
| constituents.movement.state | string | query | False | Filter by constituents.movement.state • Introduced in: 9.9 |
| constituents.movement.destination_aggregate.uuid | string | query | False | Filter by constituents.movement.destination_aggregate.uuid • Introduced in: 9.9 |
| constituents.movement.destination_aggregate.name | string | query | False | Filter by constituents.movement.destination_aggregate.name • Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---------|-------|----------|---|
| constituents.movement.percent_complete | integer | query | False | Filter by constituents.movement.percent_complete • Introduced in: 9.9 |
| constituents.name | string | query | False | Filter by constituents.name • Introduced in: 9.9 |
| access_time_enabled | boolean | query | False | Filter by access_time_enabled • Introduced in: 9.8 |
| comment | string | query | False | Filter by comment |
| snaplock.litigation_count | integer | query | False | Filter by snaplock.litigation_count |
| snaplock.unspecified_retention_file_count | integer | query | False | Filter by snaplock.unspecified_retention_file_count • Introduced in: 9.8 |
| snaplock.is_audit_log | boolean | query | False | Filter by snaplock.is_audit_log |
| snaplock.type | string | query | False | Filter by snaplock.type |
| snaplock.append_mode_enabled | boolean | query | False | Filter by snaplock.append_mode_enabled |

| Name | Type | In | Required | Description |
|----------------------------------|--------|-------|----------|---|
| snaplock.compliance_clock_time | string | query | False | Filter by snaplock.compliance_clock_time |
| snaplock.expiry_time | string | query | False | Filter by snaplock.expiry_time |
| snaplock.privileged_delete | string | query | False | Filter by snaplock.privileged_delete |
| snaplock.retention.maximum | string | query | False | Filter by snaplock.retention.maximum |
| snaplock.retention.default | string | query | False | Filter by snaplock.retention.default |
| snaplock.retention.minimum | string | query | False | Filter by snaplock.retention.minimum |
| snaplock.autocommit_period | string | query | False | Filter by snaplock.autocommit_period |
| scheduled_snapshot_naming_scheme | string | query | False | Filter by scheduled_snapshot_naming_scheme • Introduced in: 9.10 |
| style | string | query | False | Filter by style |
| type | string | query | False | Filter by type |
| create_time | string | query | False | Filter by create_time |
| clone.parent_svm.uid | string | query | False | Filter by clone.parent_svm.uid |

| Name | Type | In | Required | Description |
|------------------------------|---------|-------|----------|--|
| clone.parent_svm.name | string | query | False | Filter by clone.parent_svm.name |
| clone.parent_snapshot.name | string | query | False | Filter by clone.parent_snapshot.name |
| clone.parent_snapshot.uuid | string | query | False | Filter by clone.parent_snapshot.uuid |
| clone.split_initiated | boolean | query | False | Filter by clone.split_initiated |
| clone.split_estimate | integer | query | False | Filter by clone.split_estimate |
| clone.parent_volume.uuid | string | query | False | Filter by clone.parent_volume.uuid |
| clone.parent_volume.name | string | query | False | Filter by clone.parent_volume.name |
| clone.is_flexclone | boolean | query | False | Filter by clone.is_flexclone |
| clone.split_complete_percent | integer | query | False | Filter by clone.split_complete_percent |
| guarantee.honored | boolean | query | False | Filter by guarantee.honored |
| guarantee.type | string | query | False | Filter by guarantee.type |
| convert_unicode | boolean | query | False | Filter by convert_unicode • Introduced in: 9.10 |
| state | string | query | False | Filter by state |

| Name | Type | In | Required | Description |
|------------------------------|---------|-------|----------|--|
| consistency_group.name | string | query | False | Filter by consistency_group.name • Introduced in: 9.7 |
| consistency_group.uid | string | query | False | Filter by consistency_group.uid • Introduced in: 9.10 |
| statistics.status | string | query | False | Filter by statistics.status |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total |
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write |
| 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.latency_raw.total | integer | query | False | Filter by statistics.latency_raw.total |
| statistics.latency_raw.write | integer | query | False | Filter by statistics.latency_raw.write |
| statistics.latency_raw.other | integer | query | False | Filter by statistics.latency_raw.other |
| statistics.latency_raw.read | integer | query | False | Filter by statistics.latency_raw.read |

| Name | Type | In | Required | Description |
|------------------------------------|---------|-------|----------|--|
| statistics.throughput_raw.total | integer | query | False | Filter by statistics.throughput_raw.total |
| statistics.throughput_raw.write | integer | query | False | Filter by statistics.throughput_raw.write |
| 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.cloud.latency_raw.total | integer | query | False | Filter by statistics.cloud.latency_raw.total • Introduced in: 9.7 |
| statistics.cloud.latency_raw.write | integer | query | False | Filter by statistics.cloud.latency_raw.write • Introduced in: 9.7 |
| statistics.cloud.latency_raw.other | integer | query | False | Filter by statistics.cloud.latency_raw.other • Introduced in: 9.7 |
| statistics.cloud.latency_raw.read | integer | query | False | Filter by statistics.cloud.latency_raw.read • Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|------------------------------------|---------|-------|----------|--|
| statistics.cloud.iops_raw.total | integer | query | False | Filter by statistics.cloud.iops_raw.total • Introduced in: 9.7 |
| statistics.cloud.iops_raw.write | integer | query | False | Filter by statistics.cloud.iops_raw.write • Introduced in: 9.7 |
| statistics.cloud.iops_raw.other | integer | query | False | Filter by statistics.cloud.iops_raw.other • Introduced in: 9.7 |
| statistics.cloud.iops_raw.read | integer | query | False | Filter by statistics.cloud.iops_raw.read • Introduced in: 9.7 |
| statistics.cloud.status | string | query | False | Filter by statistics.cloud.status • Introduced in: 9.7 |
| statistics.cloud.timestamp | string | query | False | Filter by statistics.cloud.timestamp • Introduced in: 9.7 |
| statistics.flexcache_raw.timestamp | string | query | False | Filter by statistics.flexcache_raw.timestamp • Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|--|---------|-------|----------|--|
| statistics.flexcache_r aw.client_requested _blocks | integer | query | False | Filter by statistics.flexcache_r aw.client_requested _blocks • Introduced in: 9.8 |
| statistics.flexcache_r aw.cache_miss_bloc ks | integer | query | False | Filter by statistics.flexcache_r aw.cache_miss_bloc ks • Introduced in: 9.8 |
| statistics.flexcache_r aw.status | string | query | False | Filter by statistics.flexcache_r aw.status • Introduced in: 9.8 |
| statistics.timestamp | string | query | False | Filter by statistics.timestamp |
| queue_for_encryptio n | boolean | query | False | Filter by queue_for_encryptio n • Introduced in: 9.8 |
| autosize.grow_thres hold | integer | query | False | Filter by autosize.grow_thres hold |
| autosize.shrink_thre shold | integer | query | False | Filter by autosize.shrink_thre shold |
| autosize.minimum | integer | query | False | Filter by autosize.minimum |
| autosize.mode | string | query | False | Filter by autosize.mode |

| Name | Type | In | Required | Description |
|----------------------------------|---------|-------|----------|--|
| autosize.maximum | integer | query | False | Filter by autosize.maximum |
| snapmirror.destinations.is_cloud | boolean | query | False | Filter by snapmirror.destinations.is_cloud • Introduced in: 9.9 |
| snapmirror.destinations.is_ontap | boolean | query | False | Filter by snapmirror.destinations.is_ontap • Introduced in: 9.9 |
| snapmirror.is_protected | boolean | query | False | Filter by snapmirror.is_protected • Introduced in: 9.7 |
| snapshot_count | integer | query | False | Filter by snapshot_count • Introduced in: 9.10 |
| language | string | query | False | Filter by language |
| metric.iops.total | integer | query | False | Filter by metric.iops.total |
| metric.iops.write | integer | query | False | Filter by metric.iops.write |
| metric.iops.other | integer | query | False | Filter by metric.iops.other |
| metric.iops.read | integer | query | False | Filter by metric.iops.read |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total |

| Name | Type | In | Required | Description |
|-------------------------|---------|-------|----------|---|
| metric.throughput.write | integer | query | False | Filter by metric.throughput.write |
| metric.throughput.other | integer | query | False | Filter by metric.throughput.other |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read |
| metric.duration | string | query | False | Filter by metric.duration |
| metric.cloud.iops.total | integer | query | False | Filter by metric.cloud.iops.total • Introduced in: 9.7 |
| metric.cloud.iops.write | integer | query | False | Filter by metric.cloud.iops.write • Introduced in: 9.7 |
| metric.cloud.iops.other | integer | query | False | Filter by metric.cloud.iops.other • Introduced in: 9.7 |
| metric.cloud.iops.read | integer | query | False | Filter by metric.cloud.iops.read • Introduced in: 9.7 |
| metric.cloud.status | string | query | False | Filter by metric.cloud.status • Introduced in: 9.7 |

| Name | Type | In | Required | Description |
|----------------------------|---------|-------|----------|--|
| metric.cloud.duration | string | query | False | Filter by metric.cloud.duration • Introduced in: 9.7 |
| metric.cloud.latency.total | integer | query | False | Filter by metric.cloud.latency.total • Introduced in: 9.7 |
| metric.cloud.latency.write | integer | query | False | Filter by metric.cloud.latency.write • Introduced in: 9.7 |
| metric.cloud.latency.other | integer | query | False | Filter by metric.cloud.latency.other • Introduced in: 9.7 |
| metric.cloud.latency.read | integer | query | False | Filter by metric.cloud.latency.read • Introduced in: 9.7 |
| metric.cloud.timestamp | string | query | False | Filter by metric.cloud.timestamp • Introduced in: 9.7 |
| metric.status | string | query | False | Filter by metric.status |

| Name | Type | In | Required | Description |
|-------------------------------------|---------|-------|----------|---|
| metric.flexcache.timestamp | string | query | False | Filter by metric.flexcache.timestamp • Introduced in: 9.8 |
| metric.flexcache.duration | string | query | False | Filter by metric.flexcache.duration • Introduced in: 9.8 |
| metric.flexcache.cache_miss_percent | integer | query | False | Filter by metric.flexcache.cache_miss_percent • Introduced in: 9.8 |
| metric.flexcache.status | string | query | False | Filter by metric.flexcache.status • Introduced in: 9.8 |
| metric.flexcache.bandwidth_savings | integer | query | False | Filter by metric.flexcache.bandwidth_savings • Introduced in: 9.9 |
| metric.latency.total | integer | query | False | Filter by metric.latency.total |
| metric.latency.write | integer | query | False | Filter by metric.latency.write |
| metric.latency.other | integer | query | False | Filter by metric.latency.other |
| metric.latency.read | integer | query | False | Filter by metric.latency.read |

| Name | Type | In | Required | Description |
|-------------------------------------|---------|-------|----------|--|
| metric.timestamp | string | query | False | Filter by metric.timestamp |
| files.maximum | integer | query | False | Filter by files.maximum |
| files.used | integer | query | False | Filter by files.used |
| size | integer | query | False | Filter by size |
| name | string | query | False | Filter by name |
| flash_pool.caching_policy | string | query | False | Filter by flash_pool.caching_policy • Introduced in: 9.10 |
| flash_pool.cache_eligibility | string | query | False | Filter by flash_pool.cache_eligibility • Introduced in: 9.10 |
| flash_pool.cache_retention_priority | string | query | False | Filter by flash_pool.cache_retention_priority • Introduced in: 9.10 |
| application.uuid | string | query | False | Filter by application.uuid |
| application.name | string | query | False | Filter by application.name |
| 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 |
|---------------------------|---------|-------|----------|--|
| tiering.object_tags | string | query | False | Filter by tiering.object_tags • Introduced in: 9.8 |
| tiering.policy | string | query | False | Filter by tiering.policy |
| tiering.min_cooling_days | integer | query | False | Filter by tiering.min_cooling_days • Introduced in: 9.8 |
| encryption.type | string | query | False | Filter by encryption.type |
| encryption.enabled | boolean | query | False | Filter by encryption.enabled |
| encryption.key_id | string | query | False | Filter by encryption.key_id |
| encryption.state | string | query | False | Filter by encryption.state |
| 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 |
| flexgroup.uuid | string | query | False | Filter by flexgroup.uuid • Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|------------------------------------|---------|-------|----------|---|
| flexgroup.name | string | query | False | Filter by flexgroup.name • Introduced in: 9.10 |
| space.physical_used | integer | query | False | Filter by space.physical_used • Introduced in: 9.10 |
| space.snapshot.space_used_percent | integer | query | False | Filter by space.snapshot.space_used_percent • Introduced in: 9.9 |
| space.snapshot.reserve_size | integer | query | False | Filter by space.snapshot.reserve_size • Introduced in: 9.9 |
| 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.snapshot.auto_delete_trigger | string | query | False | Filter by space.snapshot.auto_delete_trigger • Introduced in: 9.10 |
| space.snapshot.reserve_available | integer | query | False | Filter by space.snapshot.reserve_available • Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---------|-------|----------|--|
| space.full_threshold_percent | integer | query | False | Filter by space.full_threshold_percent • Introduced in: 9.9 |
| space.filesystem_size | integer | query | False | Filter by space.filesystem_size • Introduced in: 9.10 |
| space.available | integer | query | False | Filter by space.available |
| space.used_by_afs | integer | query | False | Filter by space.used_by_afs • Introduced in: 9.9 |
| space.physical_used_percent | integer | query | False | Filter by space.physical_used_percent • Introduced in: 9.10 |
| space.user_data | integer | query | False | Filter by space.user_data • Introduced in: 9.10 |
| space.cross_volume_dedupe_metafiles_footprint | integer | query | False | Filter by space.cross_volume_dedupe_metafiles_footprint • Introduced in: 9.10 |
| space.metadata | integer | query | False | Filter by space.metadata |

| Name | Type | In | Required | Description |
|--|---------|-------|----------|---|
| space.block_storage_inactive_user_data | integer | query | False | Filter by space.block_storage_inactive_user_data |
| space.size_available_for_snapshots | integer | query | False | Filter by space.size_available_for_snapshots • Introduced in: 9.9 |
| space.percent_used | integer | query | False | Filter by space.percent_used • Introduced in: 9.9 |
| space.performance_tier_footprint | integer | query | False | Filter by space.performance_tier_footprint • Introduced in: 9.8 |
| space.overwrite_reserve | integer | query | False | Filter by space.overwrite_reserve • Introduced in: 9.9 |
| space.total_footprint | integer | query | False | Filter by space.total_footprint • Introduced in: 9.8 |
| space.over_provisioned | integer | query | False | Filter by space.over_provisioned |
| space.dedupe_metafiles_temporary_footprint | integer | query | False | Filter by space.dedupe_metafiles_temporary_footprint • Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|--|---------|-------|----------|--|
| space.afs_total | integer | query | False | Filter by space.afs_total • Introduced in: 9.9 |
| space.local_tier_footprint | integer | query | False | Filter by space.local_tier_footprint • Introduced in: 9.8 |
| space.block_storage_inactive_user_data_percent | integer | query | False | Filter by space.block_storage_inactive_user_data_percent • Introduced in: 9.9 |
| space.logical_space.available | integer | query | False | Filter by space.logical_space.available |
| space.logical_space.used_by_snapshots | integer | query | False | Filter by space.logical_space.used_by_snapshots • Introduced in: 9.10 |
| space.logical_space.used_by_afs | integer | query | False | Filter by space.logical_space.used_by_afs |
| space.logical_space.enforcement | boolean | query | False | Filter by space.logical_space.enforcement |
| space.logical_space.used_percent | integer | query | False | Filter by space.logical_space.used_percent • Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|---|---------|-------|----------|--|
| space.logical_space.reporting | boolean | query | False | Filter by space.logical_space.reporting |
| space.logical_space.used | integer | query | False | Filter by space.logical_space.used • Introduced in: 9.9 |
| space.used | integer | query | False | Filter by space.used |
| space.delayed_free_footprint | integer | query | False | Filter by space.delayed_free_footprint • Introduced in: 9.10 |
| space.snapshot_spill | integer | query | False | Filter by space.snapshot_spill • Introduced in: 9.10 |
| space.fractional_reserve | integer | query | False | Filter by space.fractional_reserve • Introduced in: 9.9 |
| space.cross_volume_dedupe_metafiles_temporary_footprint | integer | query | False | Filter by space.cross_volume_dedupe_metafiles_temporary_footprint • Introduced in: 9.10 |
| space.size | integer | query | False | Filter by space.size |

| Name | Type | In | Required | Description |
|--|---------|-------|----------|---|
| space.nearly_full_threshold_percent | integer | query | False | Filter by space.nearly_full_threshold_percent • Introduced in: 9.9 |
| space.snapmirror_destination_footprint | integer | query | False | Filter by space.snapmirror_destination_footprint • Introduced in: 9.10 |
| space.footprint | integer | query | False | Filter by space.footprint |
| space.filesystem_size_fixed | boolean | query | False | Filter by space.filesystem_size_fixed • Introduced in: 9.10 |
| space.dedupe_metafiles_footprint | integer | query | False | Filter by space.dedupe_metafiles_footprint • Introduced in: 9.10 |
| space.available_percent | integer | query | False | Filter by space.available_percent • Introduced in: 9.9 |
| space.overwrite_reserve_used | integer | query | False | Filter by space.overwrite_reserve_used • Introduced in: 9.9 |

| Name | Type | In | Required | Description |
|----------------------------------|---------|-------|----------|---|
| space.snapshot_reserve_unusable | integer | query | False | Filter by space.snapshot_reserve_unusable • Introduced in: 9.10 |
| space.volume_guarantee_footprint | integer | query | False | Filter by space.volume_guarantee_footprint • Introduced in: 9.10 |
| space.capacity_tier_footprint | integer | query | False | Filter by space.capacity_tier_footprint |
| space.expected_available | integer | query | False | Filter by space.expected_available • Introduced in: 9.10 |
| space.file_operation_metadata | integer | query | False | Filter by space.file_operation_metadata • Introduced in: 9.10 |
| quota.state | string | query | False | Filter by quota.state |
| aggregates.uuid | string | query | False | Filter by aggregates.uuid |
| aggregates.name | string | query | False | Filter by aggregates.name |
| is_object_store | boolean | query | False | Filter by is_object_store • Introduced in: 9.8 |

| Name | Type | In | Required | Description |
|--------------------------------|---------|-------|----------|--|
| efficiency.last_op_size | integer | query | False | Filter by efficiency.last_op_size • Introduced in: 9.9 |
| efficiency.cross_volume_dedupe | string | query | False | Filter by efficiency.cross_volume_dedupe |
| efficiency.compaction | string | query | False | Filter by efficiency.compaction |
| efficiency.op_state | string | query | False | Filter by efficiency.op_state • Introduced in: 9.9 |
| efficiency.schedule | string | query | False | Filter by efficiency.schedule • Introduced in: 9.8 |
| efficiency.last_op_state | string | query | False | Filter by efficiency.last_op_state • Introduced in: 9.9 |
| efficiency.policy.name | string | query | False | Filter by efficiency.policy.name • Introduced in: 9.7 |
| efficiency.dedupe | string | query | False | Filter by efficiency.dedupe |

| Name | Type | In | Required | Description |
|------------------------------|--------|-------|----------|--|
| efficiency.last_op_er r | string | query | False | Filter by efficiency.last_op_er r • Introduced in: 9.9 |
| efficiency.state | string | query | False | Filter by efficiency.state • Introduced in: 9.9 |
| efficiency.last_op_en d | string | query | False | Filter by efficiency.last_op_e nd • Introduced in: 9.9 |
| efficiency.last_op_be gin | string | query | False | Filter by efficiency.last_op_b egin • Introduced in: 9.9 |
| efficiency.path | string | query | False | Filter by efficiency.path • Introduced in: 9.9 |
| efficiency.progress | string | query | False | Filter by efficiency.progress • Introduced in: 9.9 |
| efficiency.type | string | query | False | Filter by efficiency.type • Introduced in: 9.9 |
| efficiency.compressi on | string | query | False | Filter by efficiency.compressi on |

| Name | Type | In | Required | Description |
|--------------------------------------|---------|-------|----------|--|
| efficiency.storage_efficiency_mode | string | query | False | Filter by efficiency.storage_efficiency_mode • Introduced in: 9.10 |
| svm.uuid | string | query | False | Filter by svm.uuid |
| svm.name | string | query | False | Filter by svm.name |
| analytics.state | string | query | False | Filter by analytics.state • Introduced in: 9.8 |
| analytics.supported | boolean | query | False | Filter by analytics.supported • Introduced in: 9.8 |
| analytics.scan_progress | integer | query | False | Filter by analytics.scan_progress • Introduced in: 9.8 |
| analytics.unsupported_reason.code | string | query | False | Filter by analytics.unsupported_reason.code • Introduced in: 9.8 |
| analytics.unsupported_reason.message | string | query | False | Filter by analytics.unsupported_reason.message • Introduced in: 9.8 |
| max_dir_size | integer | query | False | Filter by max_dir_size • Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|--------------------------------|---------|-------|----------|--|
| 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_iops | integer | query | False | Filter by qos.policy.max_throughput_iops |
| qos.policy.min_throughput_mbps | integer | query | False | Filter by qos.policy.min_throughput_mbps • Introduced in: 9.8 |
| qos.policy.name | string | query | False | Filter by qos.policy.name |
| qos.policy.max_throughput_mbps | integer | query | False | Filter by qos.policy.max_throughput_mbps |
| cloud_retrieval_policy | string | query | False | Filter by cloud_retrieval_policy • Introduced in: 9.8 |
| anti_ransomware_state | string | query | False | Filter by anti_ransomware_state • Introduced in: 9.10 |
| activity_tracking_supported | boolean | query | False | Filter by activity_tracking_supported • Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|--|---------|-------|----------|---|
| activity_tracking.state | string | query | False | Filter by activity_tracking.state • Introduced in: 9.10 |
| activity_tracking.unsupported_reason.code | string | query | False | Filter by activity_tracking.unsupported_reason.code • Introduced in: 9.10 |
| activity_tracking.unsupported_reason.message | string | query | False | Filter by activity_tracking.unsupported_reason.message • Introduced in: 9.10 |
| anti_ransomware.suspect_files.count | integer | query | False | Filter by anti_ransomware.suspect_files.count • Introduced in: 9.10 |
| anti_ransomware.suspect_files.format | string | query | False | Filter by anti_ransomware.suspect_files.format • Introduced in: 9.10 |
| anti_ransomware.attack_probability | string | query | False | Filter by anti_ransomware.attack_probability • Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|---|---------|-------|----------|--|
| anti_ransomware.state | string | query | False | Filter by anti_ransomware.state • Introduced in: 9.10 |
| anti_ransomware.dry_run_start_time | string | query | False | Filter by anti_ransomware.dry_run_start_time • Introduced in: 9.10 |
| anti_ransomware.space.used_by_logs | integer | query | False | Filter by anti_ransomware.space.used_by_logs • Introduced in: 9.10 |
| anti_ransomware.space.snapshot_count | integer | query | False | Filter by anti_ransomware.space.snapshot_count • Introduced in: 9.10 |
| anti_ransomware.space.used | integer | query | False | Filter by anti_ransomware.space.used • Introduced in: 9.10 |
| anti_ransomware.space.used_by_snapshots | integer | query | False | Filter by anti_ransomware.space.used_by_snapshots • Introduced in: 9.10 |
| anti_ransomware.attack_reports.time | string | query | False | Filter by anti_ransomware.attack_reports.time • Introduced in: 9.10 |

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

Response

Status: 200, Ok

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "activity_tracking": {
      "state": "off",
      "unsupported_reason": {
        "code": "124518405",
        "message": "Volume activity tracking cannot be enabled on
volumes that contain LUNs."
      }
    },
    "aggregates": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "analytics": {
      "scan_progress": 17,
      "state": "unknown",
      "unsupported_reason": {
        "code": "111411207",
        "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
      }
    },
    "anti_ransomware": {
      "attack_probability": "none",
      "attack_reports": {
```

```
  "_links": {
    "suspects": {
      "href": "/api/resourcelink"
    }
  },
  "time": "2021-06-01T20:36:41+05:30"
},
"dry_run_start_time": "string",
"space": {
  "snapshot_count": 0,
  "used": 0,
  "used_by_logs": 0,
  "used_by_snapshots": 0
},
"state": "disabled",
"suspect_files": {
  "count": 0,
  "format": "string"
}
},
"anti_ransomware_state": "disabled",
"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
  },
  "cloud_retrieval_policy": "default",
  "comment": "string",
  "consistency_group": {
    "name": "consistency_group_1",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "constituents": {
    "aggregates": {
      "name": "string",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "movement": {
      "cutover_window": 30,
      "destination_aggregate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "percent_complete": 0,
      "state": "replicating",
      "tiering_policy": "all"
    },
    "name": "string",
    "space": {
      "available": 0,
      "block_storage_inactive_user_data": 0,
      "capacity_tier_footprint": 0,
      "footprint": 0,
      "local_tier_footprint": 0,

```



```

    "logical_space": {
      "available": 0,
      "used_by_afs": 0
    },
    "metadata": 0,
    "over_provisioned": 0,
    "performance_tier_footprint": 0,
    "snapshot": {
      "used": 0
    },
    "total_footprint": 0,
    "used": 0
  }
},
"create_time": "2018-06-04T19:00:00Z",
"efficiency": {
  "application_io_size": "8k",
  "compaction": "inline",
  "compression": "inline",
  "cross_volume_dedupe": "inline",
  "dedupe": "inline",
  "last_op_begin": "string",
  "last_op_end": "string",
  "last_op_err": "string",
  "last_op_size": 0,
  "last_op_state": "string",
  "op_state": "idle",
  "path": "string",
  "progress": "string",
  "schedule": "string",
  "state": "disabled",
  "storage_efficiency_mode": "default",
  "type": "regular"
},
"encryption": {
  "key_id": "string",
  "key_manager_attribute": "CRN=v1:bluemix:public:containers-
kubernetes:us-south:a/asdfghjkl1234:asdfghjkl1234:worker:kubernetes-
asdfghjkl-worker1",
  "state": "encrypted",
  "status": {
    "code": "string",
    "message": "string"
  },
  "type": "none"
},

```

```

"files": {
  "used": 0
},
"flash_pool": {
  "cache_eligibility": "read",
  "cache_retention_priority": "normal",
  "caching_policy": "none"
},
"flexcache_endpoint_type": "none",
"flexgroup": {
  "name": "my_flexgroup",
  "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
},
"guarantee": {
  "type": "volume"
},
"language": "ar",
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"cloud": {
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "timestamp": "2017-01-25T11:20:13Z"
},
"duration": "PT15S",
"flexcache": {
  "bandwidth_savings": 4096,
  "cache_miss_percent": 20,
  "duration": "PT1D",
  "status": "ok",
  "timestamp": "2017-01-25T11:20:13Z"
},

```

```

    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": 0,
    "start_time": "2020-12-07T03:45:12-05:00",
    "state": "replicating",
    "tiering_policy": "all"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 100,
      "name": "default"
    },
    "junction_parent": {

```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "vs1_root",
    "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
  },
  "path": "/user/my_volume",
  "security_style": "mixed",
  "unix_permissions": 755
},
"qos": {
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "max_throughput_iops": 10000,
  "max_throughput_mbps": 500,
  "min_throughput_iops": 2000,
  "min_throughput_mbps": 500,
  "name": "performance",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
},
"quota": {
  "state": "corrupt"
},
"scheduled_snapshot_naming_scheme": "create_time",
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04T19:00:00Z",
  "expiry_time": "Wed Sep 5 11:02:42 GMT 2018",
  "is_audit_log": 1,
  "litigation_count": 10,
  "privileged_delete": "enabled",
  "retention": {
    "default": "P30Y",
    "maximum": "P30Y",
    "minimum": "P30Y"
  },
  "type": "enterprise",
  "unspecified_retention_file_count": 10
}

```

```

},
"snapshot_count": 0,
"snapshot_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"space": {
  "available": 0,
  "block_storage_inactive_user_data": 0,
  "block_storage_inactive_user_data_percent": 0,
  "capacity_tier_footprint": 0,
  "cross_volume_dedupe_metafiles_footprint": 0,
  "cross_volume_dedupe_metafiles_temporary_footprint": 0,
  "dedupe_metafiles_footprint": 0,
  "dedupe_metafiles_temporary_footprint": 0,
  "delayed_free_footprint": 0,
  "file_operation_metadata": 0,
  "filesystem_size": 0,
  "footprint": 0,
  "local_tier_footprint": 0,
  "logical_space": {
    "available": 0,
    "used": 0,
    "used_by_afs": 0,
    "used_by_snapshots": 0,
    "used_percent": 0
  },
  "metadata": 0,
  "over_provisioned": 0,
  "overwrite_reserve": 0,
  "overwrite_reserve_used": 0,
  "percent_used": 0,
  "performance_tier_footprint": 0,
  "size_available_for_snapshots": 0,
  "snapmirror_destination_footprint": 0,
  "snapshot": {
    "autodelete_trigger": "volume",
    "reserve_available": 0,
    "reserve_size": 0,
    "space_used_percent": 0,
    "used": 0
  }
}

```

```

    },
    "snapshot_reserve_unusable": 0,
    "snapshot_spill": 0,
    "total_footprint": 0,
    "used": 0,
    "user_data": 0,
    "volume_guarantee_footprint": 0
  },
  "state": "error",
  "statistics": {
    "cloud": {
      "iops_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "timestamp": "2017-01-25T11:20:13Z"
    },
    "flexcache_raw": {
      "cache_miss_blocks": 10,
      "client_requested_blocks": 500,
      "status": "ok",
      "timestamp": "2017-01-25T11:20:13Z"
    },
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
  },

```

```

    "timestamp": "2017-01-25T11:20:13Z"
  },
  "status": {
  },
  "style": "flexvol",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "object_tags": {
    },
    "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 | |

unsupported_reason

| Name | Type | Description |
|---------|--------|--|
| code | string | If volume activity tracking is not supported on the volume, this field provides an appropriate error code. |
| message | string | If volume activity tracking is not supported on the volume, this field provides an error message detailing why this is the case. |

activity_tracking

| Name | Type | Description |
|--------------------|------------------------------------|--|
| state | string | Activity tracking state of the volume. If this value is "on", ONTAP collects top metrics information for the volume in real time. There is a slight impact to I/O performance in order to collect this information. If this value is "off", no activity tracking information is collected or available to view. <ul style="list-style-type: none"> enum: ["off", "on"] Introduced in: 9.10 |
| supported | boolean | This field indicates whether or not volume activity tracking is supported on the volume. If volume activity tracking is not supported, the reason why is provided in the "activity_tracking.unsupported_reason" field. |
| unsupported_reason | unsupported_reason | |

aggregates

Aggregate

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

unsupported_reason

| Name | Type | Description |
|---------|--------|--|
| code | string | If file system analytics is not supported on the volume, this field provides the error code explaining why. |
| message | string | If file system analytics is not supported on the volume, this field provides the error message explaining why. |

analytics

| Name | Type | Description |
|--------------------|------------------------------------|---|
| scan_progress | integer | Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is <code>initializing</code> . |
| state | string | File system analytics state of the volume. If this value is "on", ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is "off", file system analytics information is not collected and not available to be viewed. If this value is "initializing", that means file system analytics was recently turned on, and the initialization scan to gather information all existing files and directories is currently running. If this value is 'unknown' that means there was an internal error when determining the file system analytics state for the volume. <ul style="list-style-type: none"> enum: ["unknown", "initializing", "off", "on"] Introduced in: 9.8 |
| supported | boolean | This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field. |
| unsupported_reason | unsupported_reason | |

_links

| Name | Type | Description |
|----------|----------------------|-------------|
| suspects | href | |

anti_ransomware_attack_report

| Name | Type | Description |
|--------|------------------------|---|
| _links | _links | |
| time | string | Timestamp at which ransomware attack is observed. |

space

| Name | Type | Description |
|-------------------|---------|--|
| snapshot_count | integer | Total number of Anti-ransomware backup Snapshot copies. |
| used | integer | Total space in bytes used by the Anti-ransomware feature. |
| used_by_logs | integer | Space in bytes used by the Anti-ransomware analytics logs. |
| used_by_snapshots | integer | Space in bytes used by the Anti-ransomware backup Snapshot copies. |

suspect_files

| Name | Type | Description |
|--------|---------|---|
| count | integer | Total number of <code>suspect_files.format</code> files observed by the Anti-ransomware analytics engine on the volume. |
| format | string | File formats observed by the Anti-ransomware analytics engine on the volume. |

anti_ransomware

Anti-ransomware related information of the volume.

| Name | Type | Description |
|--------------------|--|--|
| attack_probability | string | Probability of a ransomware attack. <code>none</code> No files are suspected of ransomware activity. <code>low</code> A number of files are suspected of ransomware activity. <code>moderate</code> A moderate number of files are suspected of ransomware activity. <code>high</code> A large number of files are suspected of ransomware activity. |
| attack_reports | array[anti_ransomware_attack_report] | |
| dry_run_start_time | string | Time when Anti-ransomware monitoring state is set to dry-run value for starting evaluation mode. |
| space | space | |

| Name | Type | Description |
|---------------|----------------------|--|
| state | string | <p>Anti-ransomware state.</p> <p><code>disabled</code> Anti-ransomware monitoring is disabled on the volume. This is the default state in a POST operation.</p> <p><code>disable_in_progress</code> Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. Valid in GET operation.</p> <p><code>dry_run</code> Anti-ransomware monitoring is enabled in the evaluation mode.</p> <p><code>enabled</code> Anti-ransomware monitoring is active on the volume.</p> <p><code>paused</code> Anti-ransomware monitoring is paused on the volume.</p> <p><code>enable_paused</code> Anti-ransomware monitoring is paused on the volume from its earlier enabled state. Valid in GET operation.</p> <p><code>dry_run_paused</code> Anti-ransomware monitoring is paused on the volume from its earlier <code>dry_run</code> state. Valid in GET operation. For POST, the valid Anti-ransomware states are only <code>disabled</code>, <code>enabled</code> and <code>dry_run</code>, whereas for PATCH, <code>paused</code> is also valid along with the three valid states for POST.</p> |
| suspect_files | array[suspect_files] | |

application

| Name | Type | Description |
|------|--------|--|
| name | string | Name of the application to which the volume belongs. Available only when the volume is part of an application. |
| uuid | string | UUID of the application to which the volume belongs. Available only when the volume is part of an application. |

autosize

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

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

snapshot_reference

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

parent_svm

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

parent_volume

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

| Name | Type | Description |
|------|--------|--|
| uuid | string | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6 |

clone

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

consistency_group

Consistency group the volume is part of.

| Name | Type | Description |
|------|--------|--|
| name | string | The name of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |
| uuid | string | The UUID of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |

aggregates

| Name | Type | Description |
|------|--------|--|
| name | string | Name of the aggregate hosting the FlexGroup Constituent. |
| uuid | string | Unique identifier for the aggregate. |

destination_aggregate

Aggregate

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

movement

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

| Name | Type | Description |
|-----------------------|---------------------------------------|---|
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

logical_space

| Name | Type | Description |
|-------------|---------|--|
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |

| Name | Type | Description |
|-------------|---------|---|
| 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 |
|----------------------------------|---------|---|
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| block_storage_inactive_user_data | integer | The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |

| Name | Type | Description |
|----------------------------|-------------------------------|--|
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| snapshot | snapshot | |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| used_percent | integer | The virtual space used (includes volume reserves) before storage efficiency, as a percent. |

constituents

| Name | Type | Description |
|------------|----------------------------|-------------|
| aggregates | aggregates | |

| 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 | FlexGroup Constituents name |
| space | space | |

policy

| Name | Type | Description |
|------|--------|---|
| name | string | Specifies the name of the efficiency policy. Valid for PATCH. |

efficiency

| Name | Type | Description |
|---------------------|--------|--|
| application_io_size | string | Block size to use by compression. Valid for POST. |
| 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. |

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

| 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. |
| last_op_begin | string | Last sis operation begin timestamp. |
| last_op_end | string | Last sis operation end timestamp. |
| last_op_err | string | Last sis operation error text. |
| last_op_size | integer | Last sis operation size. |
| last_op_state | string | Last sis operation state. |
| op_state | string | Sis status of the volume. |
| path | string | Absolute volume path of the volume. |
| policy | policy | |
| progress | string | Sis progress of the volume. |
| schedule | string | Schedule associated with volume. |

| Name | Type | Description |
|-------------------------|--------|--|
| state | string | Storage efficiency state of the volume. Currently, this field supports POST/PATCH only for RW (Read-Write) volumes on FSx for ONTAP. disabled ‐ All storage efficiency features are disabled. mixed ‐ Read-only field for FlexGroup volumes, storage efficiency is enabled on certain constituents and disabled on others. On FSx for ONTAP ‐ enabled ‐ All supported storage efficiency features for the volume are enabled. custom ‐ Read-only field currently only supported for the FSx for ONTAP, user-defined storage efficiency features are enabled. For other platforms ‐ enabled ‐ At least one storage efficiency feature for the volume is enabled. <ul style="list-style-type: none"> enum: ["disabled", "enabled", "mixed", "custom"] Introduced in: 9.9 |
| storage_efficiency_mode | string | Storage efficiency mode used by volume. This parameter is supported only on AFF platform. |
| type | string | Sis Type of the volume. |

status

| Name | Type | Description |
|---------|--------|-----------------------------------|
| code | string | Encryption progress message code. |
| message | string | Encryption progress message. |

encryption

| Name | Type | Description |
|-----------------------|---------|--|
| enabled | boolean | Creates an encrypted or an unencrypted volume. For POST, when set to 'true', a new key is generated and used to encrypt the given volume. In that case, the underlying SVM must be configured with the key manager. When set to 'false', the volume created will be unencrypted. For PATCH, when set to 'true', it encrypts an unencrypted volume. Specifying the parameter as 'false' in a PATCH operation for an encrypted volume is only supported when moving the volume to another aggregate. |
| key_id | string | The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key. |
| key_manager_attribute | string | Specifies an additional key manager attribute that is an identifier-value pair, separated by '='. For example, CRN=unique-value. This parameter is required when using the POST method and an IBM Key Lore key manager is configured on the SVM. |
| rekey | boolean | If set to 'true', re-encrypts the volume with a new key. Valid in PATCH. |
| state | string | Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one. |

| Name | Type | Description |
|--------|------------------------|---|
| status | status | |
| type | string | Volume encryption type. none ‐ The volume is a plain-text one. volume ‐ The volume is encrypted with NVE (NetApp Volume Encryption). aggregate ‐ The volume is encrypted with NAE (NetApp Aggregate Encryption). |

error_state

| Name | Type | Description |
|-----------------|---------|---|
| has_bad_blocks | boolean | Indicates whether the volume has any corrupt data blocks. If the damaged data block is accessed, an IO error, such as EIO for NFS or STATUS_FILE_CORRUPT for CIFS, is returned. |
| is_inconsistent | boolean | Indicates whether the file system has any inconsistencies. true ‐ File system is inconsistent. false ‐ File system 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. |

flash_pool

| Name | Type | Description |
|--------------------------|--------|---|
| cache_eligibility | string | If this parameter is specified, the command displays information only about the volume or volumes with the specified Flash Pool caching attributes. |
| cache_retention_priority | string | If this parameter is specified, the command displays the volumes that match the specified cache retention priority policy. A cache retention priority defines how long the blocks of a volume will be cached in the Flash Pool once they become cold. |

| Name | Type | Description |
|----------------|--------|---|
| caching_policy | string | This optionally specifies the caching policy to apply to the volume. A caching policy defines how the system caches a volume's data in Flash Cache modules. If a caching policy is not assigned to a volume, the system uses the caching policy that is assigned to the containing SVM. If a caching policy is not assigned to the containing SVM, the system uses the default cluster-wide policy. |

flexgroup

| Name | Type | Description |
|------|--------|---|
| name | string | Name of the FlexGroup volume that the constituent is part of. |
| uuid | string | Unique identifier for the FlexGroup volume that the constituent is part of. |

guarantee

| Name | Type | Description |
|---------|---------|---|
| honored | boolean | Is the space guarantee of this volume honored in the aggregate? |
| type | string | The type of space guarantee of this volume in the aggregate. |

iops

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

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

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

latency

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

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

cloud

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.

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

| Name | Type | Description |
|-----------|---------|---|
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

flexcache

Performance number for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
|--------------------|---------|--|
| bandwidth_savings | integer | Bandwidth savings denoting the amount of data served locally by the cache, in bytes. |
| cache_miss_percent | integer | Cache miss percentage. |

| Name | Type | Description |
|-----------|--------|---|
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

throughput

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

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

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

metric

Performance numbers, such as IOPS, latency and throughput.

| Name | Type | Description |
|-----------|---------------------------|--|
| _links | _links | |
| cloud | cloud | Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| flexcache | flexcache | Performance number for FlexCache used to measure cache effectiveness. |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
|------------|----------------------------|---|
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

movement

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

| Name | Type | Description |
|-----------------------|---------------------------------------|---|
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| start_time | string | Start time of volume move. |

| Name | Type | Description |
|----------------|--------|---|
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

export_policy

Export Policy

| Name | Type | Description |
|--------|------------------------|-------------|
| _links | _links | |
| id | integer | |
| name | string | |

junction_parent

| Name | Type | Description |
|--------|------------------------|--|
| _links | _links | |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |
| uuid | string | Unique identifier for the parent volume. |

nas

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

| Name | Type | Description |
|------------------|---------|---|
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth digit), first digit is assumed to be zero. |

policy

| Name | Type | Description |
|------------------------|------------------------|---|
| _links | _links | |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. |

| Name | Type | Description |
|---------------------|---------|---|
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

QoS information

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

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

| Name | Type | Description |
|---------|---------|--|
| enabled | boolean | This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVols or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVols or FlexGroup volumes when the quota state is "off". |
| state | string | Quota state of the volume |

retention

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

destinations

| Name | Type | Description |
|----------|---------|---|
| is_cloud | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to a cloud destination. |
| is_ontap | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to an ONTAP destination. <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.9 |

snapmirror

Specifies attributes for SnapMirror protection.

| Name | Type | Description |
|--------------|------------------------------|---|
| destinations | destinations | |
| is_protected | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data. |

snapshot_policy

This is a reference to the Snapshot copy policy.

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

logical_space

| Name | Type | Description |
|-------------|---------|---|
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |
| used | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), in bytes. |
| used_by_afs | integer | The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes. |

| Name | Type | Description |
|-------------------|---------|--|
| used_by_snapshots | integer | Size that is logically used across all Snapshot copies in the volume, in bytes. |
| used_percent | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), as a percentage. |

snapshot

| Name | Type | Description |
|--------------------|---------|---|
| autodelete_enabled | boolean | Specifies whether Snapshot copy autodelete is currently enabled on this volume. |
| autodelete_trigger | string | Specifies when the system should trigger an autodelete of Snapshot copies. When set to <i>volume</i> , autodelete is triggered based on volume fullness. When set to <i>snap_reserve</i> , autodelete is triggered based on Snapshot reserve fullness. The default value is <i>volume</i> . |
| reserve_available | integer | Size available for Snapshot copies within the Snapshot copy reserve, in bytes. |
| reserve_percent | integer | The space that has been set aside as a reserve for Snapshot copy usage, in percent. |
| reserve_size | integer | Size in the volume that has been set aside as a reserve for Snapshot copy usage, in bytes. |
| space_used_percent | integer | Percentage of snapshot reserve size that has been used. |
| used | integer | The total space used by Snapshot copies in the volume, in bytes. |

space

| Name | Type | Description |
|---|---------|---|
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| 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. |
| block_storage_inactive_user_data_percent | integer | Percentage of size that is physically used in the performance tier of the volume. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |
| cross_volume_dedupe_metafiles_footprint | integer | Cross volume deduplication metadata footprint, in bytes. |
| cross_volume_dedupe_metafiles_temporary_footprint | integer | Cross volume temporary deduplication metadata footprint, in bytes. |
| dedupe_metafiles_footprint | integer | Deduplication metadata footprint, in bytes. |
| dedupe_metafiles_temporary_footprint | integer | Temporary deduplication metadata footprint, in bytes. |
| delayed_free_footprint | integer | Delayed free blocks footprint, in bytes. |
| expected_available | integer | Size that should be available for the volume, irrespective of available size in the aggregate, in bytes. |

| Name | Type | Description |
|-------------------------------|-------------------------------|---|
| file_operation_metadata | integer | File operation metadata footprint, in bytes. |
| filesystem_size | integer | Total usable size of the volume, in bytes. |
| filesystem_size_fixed | boolean | Specifies whether the file system is to remain of the same size when set to true or to grow when set to false. This option is automatically set to true when a volume becomes SnapMirrored. |
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| fractional_reserve | integer | Used to change the amount of space reserved for overwrites of reserved objects in a volume. |
| full_threshold_percent | integer | Volume full threshold percentage at which EMS warnings can be sent. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| nearly_full_threshold_percent | integer | Volume nearly full threshold percentage at which EMS warnings can be sent. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| overwrite_reserve | integer | Reserved space for overwrites, in bytes. |
| overwrite_reserve_used | integer | Overwrite logical reserve space used, in bytes. |

| Name | Type | Description |
|----------------------------------|--------------------------|---|
| percent_used | integer | Percentage of the volume size that is used. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| physical_used | integer | Size that is physically used in the volume, in bytes. |
| physical_used_percent | integer | Size that is physically used in the volume, as a percentage. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| size_available_for_snapshots | integer | Available space for Snapshot copies from snap-reserve, in bytes. |
| snapmirror_destination_footprint | integer | SnapMirror destination footprint, in bytes. |
| snapshot | snapshot | |
| snapshot_reserve_unusable | integer | Snapshot reserve that is not available for Snapshot copy creation, in bytes. |
| snapshot_spill | integer | Space used by the snapshot copies beyond the snap-reserve, in bytes. |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| user_data | integer | User data, in bytes. |

| Name | Type | Description |
|----------------------------|---------|---|
| volume_guarantee_footprint | integer | Space reserved for future writes in the volume, in bytes. |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

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

latency_raw

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

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

cloud

These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.

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

flexcache_raw

Performance numbers for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
|-------------------------|---------|---|
| cache_miss_blocks | integer | Blocks retrieved from origin in case of a cache miss. This can be divided by the raw client_requested_blocks and multiplied by 100 to calculate the cache miss percentage. |
| client_requested_blocks | integer | Total blocks requested by the client. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

throughput_raw

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

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

statistics

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

| Name | Type | Description |
|---------------|-------------------------------|--|
| cloud | cloud | These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools. |
| flexcache_raw | flexcache_raw | Performance numbers for FlexCache used to measure cache effectiveness. |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |

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

svm

SVM containing the volume. Required on POST.

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

tiering

| Name | Type | Description |
|------------------|---------------|--|
| min_cooling_days | integer | <p>This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. Note that this parameter is only used for tiering purposes and does not affect the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data. This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy.</p> |
| object_tags | array[string] | <p>This parameter specifies tags of a volume for objects stored on a FabricPool-enabled aggregate. Each tag is a key,value pair and should be in the format "key=value".</p> |

| Name | Type | Description |
|--------|--------|--|
| policy | string | <p>Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.</p> <p>all &dash; This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.</p> <p>auto &dash; This policy allows tiering of both snapshot and active file system user data to the cloud store</p> <p>none &dash; Volume blocks will not be tiered to the cloud store.</p> <p>snapshot_only &dash; This policy allows tiering of only the volume Snapshot copies not associated with the active file system.</p> <p>The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days.</p> |

| Name | Type | Description |
|-----------|---------|---|
| supported | boolean | This parameter specifies whether or not FabricPools are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only FabricPool aggregates are used if this parameter is set to true and only non FabricPool aggregates are used if this parameter is set to false. Tiering support for a FlexGroup can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only does the volume need to support tiering by using FabricPools, the tiering "policy" must not be 'none'. A volume that uses FabricPools but has a tiering "policy" of 'none' supports tiering, but will not tier any data. |

volume

| Name | Type | Description |
|---------------------|-------------------------------------|---|
| _links | _links | |
| access_time_enabled | boolean | Indicates whether or not access time updates are enabled on the volume. |
| activity_tracking | activity_tracking | |
| aggregates | array[aggregates] | Aggregate hosting the volume. Required on POST. |
| analytics | analytics | |
| anti_ransomware | anti_ransomware | Anti-ransomware related information of the volume. |

| Name | Type | Description |
|-----------------------|-----------------------------|---|
| anti_ransomware_state | string | <p>The Anti-ransomware state of the volume. If no "anti_ransomware_state" property is specified, the volume inherits the value from its parent SVM's "anti_ransomware_default_volume_state" property. If this value is "disabled", Anti-ransomware is disabled on the volume. If this value is "enabled", Anti-ransomware is enabled on the volume and alerts are raised if any suspect is detected for those volumes. If this value is "dry_run", Anti-ransomware is enabled in the dry-run or learning mode on the volume. The "dry_run" state is same as the "enabled" state except that the analytics data is used here for learning. No alerts are raised for any detections or violations. If this value is "paused", Anti-ransomware is paused on the volume. Additionally, three more states are available, which are only valid for GET. If this value is "disable_in_progress", Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. If this value is "enable_paused", Anti-ransomware is paused on the volume from its earlier enabled state. If this value is "dry_run_paused", Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. For POST, the valid Anti-ransomware states are only "disabled", "enabled" and "dry_run", whereas for PATCH, "paused" is also valid along with the three valid states for POST.</p> |
| application | application | |
| autosize | autosize | |
| clone | clone | |

| Name | Type | Description |
|------------------------|---------------------------------------|---|
| cloud_retrieval_policy | string | <p>This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved. "on_read" policy retrieves tiered data for all client driven data reads. "never" policy never retrieves tiered data. "promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved.</p> |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| consistency_group | consistency_group | Consistency group the volume is part of. |
| constituents | array[constituents] | |

| 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 volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list. |
| convert_unicode | boolean | Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients. |
| 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 | |
| flash_pool | flash_pool | |
| 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. |
| flexgroup | flexgroup | |
| guarantee | guarantee | |

| Name | Type | Description |
|-----------------|--------------------------|--|
| is_object_store | boolean | Specifies whether the volume is provisioned for an object store server. |
| is_svm_root | boolean | Specifies whether the volume is a root volume of the SVM it belongs to. |
| language | string | Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting. |
| max_dir_size | integer | Maximum directory size. This value sets maximum size, in bytes, to which a directory can grow. The default maximum directory size for FlexVol volumes is model-dependent, and optimized for the size of system memory. Before increasing the maximum directory size, involve technical support. |
| 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 | |

| Name | Type | Description |
|----------------------------------|-----------------|---|
| qos | qos | QoS information |
| queue_for_encryption | boolean | Specifies whether the volume is queued for encryption. |
| quota | quota | Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume. |
| scheduled_snapshot_naming_scheme | string | <p>Naming Scheme for automatic Snapshot copies:</p> <ul style="list-style-type: none"> • create_time - Automatic Snapshot copies are saved as per the start of their current date and time. • ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> |
| size | integer | Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size. |
| snaplock | snaplock | |
| snapmirror | snapmirror | Specifies attributes for SnapMirror protection. |
| snapshot_count | integer | Number of Snapshot copies in the volume. |
| snapshot_policy | snapshot_policy | This is a reference to the Snapshot copy policy. |
| space | space | |

| Name | Type | Description |
|------------|----------------------------|---|
| state | string | Volume state. Client access is supported only when volume is online and junctioned. Taking volume to offline or restricted state removes its junction path and blocks client access. When volume is in restricted state some operations like parity reconstruction and iron on commit are allowed. 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. |
| status | array[string] | Describes the current status of a volume. |

| Name | Type | Description |
|---------|---------|---|
| style | string | <p>The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol" you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. The style "flexgroup_constituent" is not supported when creating a volume. flexvol &dash; flexible volumes and FlexClone volumes flexgroup &dash; FlexGroup volumes flexgroup_constituent &dash; FlexGroup constituents.</p> |
| svm | svm | SVM containing the volume. Required on POST. |
| tiering | tiering | |
| type | string | <p>Type of the volume. rw &dash; read-write volume. dp &dash; data-protection volume. ls &dash; load-sharing <code>dp</code> volume. Valid in GET.</p> |

| Name | Type | Description |
|-------------------------|---------|--|
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration. |
| uuid | string | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6 |

error_arguments

| Name | Type | Description |
|---------|--------|------------------|
| code | string | 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

Introduced In: 9.6

Creates a volume on a specified SVM and storage aggregates.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the volume.
- `name` - Name of the volume.
- `aggregates.name` or `aggregates.uuid` - Existing aggregates in which to create the volume.

Default property values

- `state` - *online*
- `size` - *20MB*
- `style` - *flexvol*
- `type` - *rw*
- `encryption.enabled` - *false*
- `snapshot_policy.name` - *default*
- `guarantee.type` - *volume*
- `anti_ransomware.state` - *default*

Related ONTAP commands

- `volume create`
- `volume clone create`

Parameters

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

Request Body

| Name | Type | Description |
|---------------------|------------------------|---|
| _links | _links | |
| access_time_enabled | boolean | Indicates whether or not access time updates are enabled on the volume. |

| Name | Type | Description |
|-------------------|-------------------------------------|---|
| activity_tracking | activity_tracking | |
| aggregates | array[aggregates] | Aggregate hosting the volume. Required on POST. |
| analytics | analytics | |
| anti_ransomware | anti_ransomware | Anti-ransomware related information of the volume. |

| Name | Type | Description |
|-----------------------|-------------|---|
| anti_ransomware_state | string | <p>The Anti-ransomware state of the volume. If no "anti_ransomware_state" property is specified, the volume inherits the value from its parent SVM's "anti_ransomware_default_volume_state" property. If this value is "disabled", Anti-ransomware is disabled on the volume. If this value is "enabled", Anti-ransomware is enabled on the volume and alerts are raised if any suspect is detected for those volumes. If this value is "dry_run", Anti-ransomware is enabled in the dry-run or learning mode on the volume. The "dry_run" state is same as the "enabled" state except that the analytics data is used here for learning. No alerts are raised for any detections or violations. If this value is "paused", Anti-ransomware is paused on the volume. Additionally, three more states are available, which are only valid for GET. If this value is "disable_in_progress", Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. If this value is "enable_paused", Anti-ransomware is paused on the volume from its earlier enabled state. If this value is "dry_run_paused", Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. For POST, the valid Anti-ransomware states are only "disabled", "enabled" and "dry_run", whereas for PATCH, "paused" is also valid along with the three valid states for POST.</p> |
| application | application | |
| autosize | autosize | |
| clone | clone | |

| Name | Type | Description |
|------------------------|---------------------------------------|---|
| cloud_retrieval_policy | string | <p>This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved. "on_read" policy retrieves tiered data for all client driven data reads. "never" policy never retrieves tiered data. "promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved.</p> |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| consistency_group | consistency_group | Consistency group the volume is part of. |
| constituents | array[constituents] | |

| 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 volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list. |
| convert_unicode | boolean | Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients. |
| 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 | |
| flash_pool | flash_pool | |
| 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. |
| flexgroup | flexgroup | |
| guarantee | guarantee | |
| is_object_store | boolean | Specifies whether the volume is provisioned for an object store server. |

| Name | Type | Description |
|----------------------|--------------------------|--|
| is_svm_root | boolean | Specifies whether the volume is a root volume of the SVM it belongs to. |
| language | string | Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting. |
| max_dir_size | integer | Maximum directory size. This value sets maximum size, in bytes, to which a directory can grow. The default maximum directory size for FlexVol volumes is model-dependent, and optimized for the size of system memory. Before increasing the maximum directory size, involve technical support. |
| 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 |
| queue_for_encryption | boolean | Specifies whether the volume is queued for encryption. |

| Name | Type | Description |
|----------------------------------|---------------------------------|---|
| quota | quota | Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume. |
| scheduled_snapshot_naming_scheme | string | <p>Naming Scheme for automatic Snapshot copies:</p> <ul style="list-style-type: none"> • create_time - Automatic Snapshot copies are saved as per the start of their current date and time. • ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> |
| size | integer | Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size. |
| snaplock | snaplock | |
| snapmirror | snapmirror | Specifies attributes for SnapMirror protection. |
| snapshot_count | integer | Number of Snapshot copies in the volume. |
| snapshot_policy | snapshot_policy | This is a reference to the Snapshot copy policy. |
| space | space | |

| Name | Type | Description |
|------------|----------------------------|---|
| state | string | Volume state. Client access is supported only when volume is online and junctioned. Taking volume to offline or restricted state removes its junction path and blocks client access. When volume is in restricted state some operations like parity reconstruction and iron on commit are allowed. 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. |
| status | array[string] | Describes the current status of a volume. |

| Name | Type | Description |
|---------|---------|---|
| style | string | <p>The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol" you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. The style "flexgroup_constituent" is not supported when creating a volume. flexvol &dash; flexible volumes and FlexClone volumes flexgroup &dash; FlexGroup volumes flexgroup_constituent &dash; FlexGroup constituents.</p> |
| svm | svm | SVM containing the volume. Required on POST. |
| tiering | tiering | |
| type | string | <p>Type of the volume. rw &dash; read-write volume. dp &dash; data-protection volume. ls &dash; load-sharing <code>dp</code> volume. Valid in GET.</p> |

| 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 • Introduced in: 9.6 |

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "activity_tracking": {
    "state": "off",
    "unsupported_reason": {
      "code": "124518405",
      "message": "Volume activity tracking cannot be enabled on volumes
that contain LUNs."
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "analytics": {
    "scan_progress": 17,
    "state": "unknown",
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
    }
  },
  "anti_ransomware": {
    "attack_probability": "none",
    "attack_reports": {
      "_links": {
        "suspects": {
          "href": "/api/resourcelink"
        }
      }
    },
    "time": "2021-06-01T20:36:41+05:30"
  },
  "dry_run_start_time": "string",
  "space": {
```

```
    "snapshot_count": 0,
    "used": 0,
    "used_by_logs": 0,
    "used_by_snapshots": 0
  },
  "state": "disabled",
  "suspect_files": {
    "count": 0,
    "format": "string"
  }
},
"anti_ransomware_state": "disabled",
"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
  },
  "cloud_retrieval_policy": "default",
  "comment": "string",
  "consistency_group": {
    "name": "consistency_group_1",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "constituents": {
    "aggregates": {
      "name": "string",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "movement": {
      "cutover_window": 30,
      "destination_aggregate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "percent_complete": 0,
      "state": "replicating",
      "tiering_policy": "all"
    },
    "name": "string",
    "space": {
      "available": 0,
      "block_storage_inactive_user_data": 0,
      "capacity_tier_footprint": 0,
      "footprint": 0,
      "local_tier_footprint": 0,
      "logical_space": {
        "available": 0,
        "used_by_afs": 0
      },
      "metadata": 0,
      "over_provisioned": 0,
      "performance_tier_footprint": 0,
      "snapshot": {
        "used": 0
      }
    }
  }
}

```

```

    },
    "total_footprint": 0,
    "used": 0
  }
},
"create_time": "2018-06-04T19:00:00Z",
"efficiency": {
  "application_io_size": "8k",
  "compaction": "inline",
  "compression": "inline",
  "cross_volume_dedupe": "inline",
  "dedupe": "inline",
  "last_op_begin": "string",
  "last_op_end": "string",
  "last_op_err": "string",
  "last_op_size": 0,
  "last_op_state": "string",
  "op_state": "idle",
  "path": "string",
  "progress": "string",
  "schedule": "string",
  "state": "disabled",
  "storage_efficiency_mode": "default",
  "type": "regular"
},
"encryption": {
  "key_id": "string",
  "key_manager_attribute": "CRN=v1:bluemix:public:containers-
kubernetes:us-south:a/asdfghjkl1234:asdfghjkl1234:worker:kubernetes-
asdfghjkl-worker1",
  "state": "encrypted",
  "status": {
    "code": "string",
    "message": "string"
  },
  "type": "none"
},
"files": {
  "used": 0
},
"flash_pool": {
  "cache_eligibility": "read",
  "cache_retention_priority": "normal",
  "caching_policy": "none"
},
"flexcache_endpoint_type": "none",

```

```
"flexgroup": {
  "name": "my_flexgroup",
  "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
},
"guarantee": {
  "type": "volume"
},
"language": "ar",
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"cloud": {
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "timestamp": "2017-01-25T11:20:13Z"
},
"duration": "PT15S",
"flexcache": {
  "bandwidth_savings": 4096,
  "cache_miss_percent": 20,
  "duration": "PT1D",
  "status": "ok",
  "timestamp": "2017-01-25T11:20:13Z"
},
"iops": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"latency": {
  "read": 200,
  "total": 1000,
  "write": 100
}
```



```

    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": 0,
    "start_time": "2020-12-07T03:45:12-05:00",
    "state": "replicating",
    "tiering_policy": "all"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 100,
      "name": "default"
    },
    "junction_parent": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "vs1_root",
      "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
    },
    "path": "/user/my_volume",

```

```

    "security_style": "mixed",
    "unix_permissions": 755
  },
  "qos": {
    "policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    },
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"quota": {
  "state": "corrupt"
},
"scheduled_snapshot_naming_scheme": "create_time",
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04T19:00:00Z",
  "expiry_time": "Wed Sep 5 11:02:42 GMT 2018",
  "is_audit_log": 1,
  "litigation_count": 10,
  "privileged_delete": "enabled",
  "retention": {
    "default": "P30Y",
    "maximum": "P30Y",
    "minimum": "P30Y"
  },
  "type": "enterprise",
  "unspecified_retention_file_count": 10
},
"snapshot_count": 0,
"snapshot_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"name": "default",

```

```
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"space": {
  "available": 0,
  "block_storage_inactive_user_data": 0,
  "block_storage_inactive_user_data_percent": 0,
  "capacity_tier_footprint": 0,
  "cross_volume_dedupe_metafiles_footprint": 0,
  "cross_volume_dedupe_metafiles_temporary_footprint": 0,
  "dedupe_metafiles_footprint": 0,
  "dedupe_metafiles_temporary_footprint": 0,
  "delayed_free_footprint": 0,
  "file_operation_metadata": 0,
  "filesystem_size": 0,
  "footprint": 0,
  "local_tier_footprint": 0,
  "logical_space": {
    "available": 0,
    "used": 0,
    "used_by_afs": 0,
    "used_by_snapshots": 0,
    "used_percent": 0
  },
  "metadata": 0,
  "over_provisioned": 0,
  "overwrite_reserve": 0,
  "overwrite_reserve_used": 0,
  "percent_used": 0,
  "performance_tier_footprint": 0,
  "size_available_for_snapshots": 0,
  "snapmirror_destination_footprint": 0,
  "snapshot": {
    "autodelete_trigger": "volume",
    "reserve_available": 0,
    "reserve_size": 0,
    "space_used_percent": 0,
    "used": 0
  },
  "snapshot_reserve_unusable": 0,
  "snapshot_spill": 0,
  "total_footprint": 0,
  "used": 0,
  "user_data": 0,
  "volume_guarantee_footprint": 0
},
"state": "error",
```

```

"statistics": {
  "cloud": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "flexcache_raw": {
    "cache_miss_blocks": 10,
    "client_requested_blocks": 500,
    "status": "ok",
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "iops_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"status": {
},
"style": "flexvol",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
}

```

```

    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"tiering": {
  "object_tags": {
  },
  "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 |
|------------|---|
| 787140 | One of "aggregates.uuid", "aggregates.name", or "style" must be provided. |
| 787141 | The specified "aggregates.name" and "aggregates.uuid" refer to different aggregates. |
| 917526 | The volume name specified is a duplicate. |
| 917829 | Volume autosize grow threshold must be larger than autosize shrink threshold. |
| 917831 | Volume minimum autosize must be smaller than the maximum autosize. |
| 917835 | Maximum allowed snapshot.reserve_percent value during a volume creation is 90. Use PATCH to set it to a higher value after the volume has been created. |
| 918191 | Flexvol tiering min cooling days requires an effective cluster version of ONTAP 9.4 or later. |
| 918194 | Tiering min cooling days not supported for SVMDR. |
| 918195 | Tiering min cooling days not supported for non data volumes. |
| 918196 | Tiering min cooling days not allowed for the provided tiering policy. |
| 918215 | FlexGroup tiering min cooling days requires an effective cluster version of ONTAP 9.5 or later. |
| 918233 | The target field cannot be specified for this operation. |
| 918236 | The specified "parent_volume.uuid" and "parent_volume.name" do not refer to the same volume. |
| 918240 | The target style is an invalid volume style. |
| 918241 | The target style is an unsupported volume style for volume creation. |
| 918242 | When creating a flexible volume, exactly one aggregate must be specified via either "aggregates.name" or "aggregates.uuid". |
| 918243 | The specified Snapshot copy UUID is not correct for the specified Snapshot copy name. |
| 918244 | Invalid "volume.type" for clone volume. |
| 918246 | "volume.clone.parent_volume.name" or "volume.clone.parent_volume.uuid" must be provided. |
| 918247 | Specifying a value is not valid for a volume FlexClone creation. |
| 918252 | "nas.path" is invalid. |

| Error Code | Description |
|------------|---|
| 918290 | cloud retrieval policy requires an effective cluster version of 9.8 or later. |
| 918291 | Invalid volume cloud retrieval policy for the provided tiering policy. |
| 918292 | cloud retrieval policy not supported for non data volume. |
| 918521 | The volume maximum autosize must be smaller than or equal to the maximum volume size. |
| 918524 | Volume minimum autosize must be less than or equal to the current volume size. |
| 2621706 | The specified "svm.uuid" and "svm.name" do not refer to the same SVM. |
| 2621707 | No SVM was specified. Either "svm.name" or "svm.uuid" must be supplied. |
| 111411205 | File system analytics requires an effective cluster version of 9.8 or later. |
| 111411206 | The specified "analytics.state" is invalid. |
| 111411207 | File system analytics cannot be enabled on volumes that contain LUNs. |

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

Example error

```

{
  "error": {
    "arguments": {
      "code": "string",
      "message": "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 | |

unsupported_reason

| Name | Type | Description |
|---------|--------|--|
| code | string | If volume activity tracking is not supported on the volume, this field provides an appropriate error code. |
| message | string | If volume activity tracking is not supported on the volume, this field provides an error message detailing why this is the case. |

activity_tracking

| Name | Type | Description |
|-------|--------|--|
| state | string | <p>Activity tracking state of the volume. If this value is "on", ONTAP collects top metrics information for the volume in real time. There is a slight impact to I/O performance in order to collect this information. If this value is "off", no activity tracking information is collected or available to view.</p> <ul style="list-style-type: none">• enum: ["off", "on"]• Introduced in: 9.10 |

| Name | Type | Description |
|--------------------|------------------------------------|--|
| supported | boolean | This field indicates whether or not volume activity tracking is supported on the volume. If volume activity tracking is not supported, the reason why is provided in the "activity_tracking.unsupported_reason" field. |
| unsupported_reason | unsupported_reason | |

aggregates

Aggregate

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

unsupported_reason

| Name | Type | Description |
|---------|--------|--|
| code | string | If file system analytics is not supported on the volume, this field provides the error code explaining why. |
| message | string | If file system analytics is not supported on the volume, this field provides the error message explaining why. |

analytics

| Name | Type | Description |
|---------------|---------|---|
| scan_progress | integer | Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is <code>initializing</code> . |

| Name | Type | Description |
|--------------------|------------------------------------|---|
| state | string | File system analytics state of the volume. If this value is "on", ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is "off", file system analytics information is not collected and not available to be viewed. If this value is "initializing", that means file system analytics was recently turned on, and the initialization scan to gather information all existing files and directories is currently running. If this value is 'unknown' that means there was an internal error when determining the file system analytics state for the volume. <ul style="list-style-type: none"> enum: ["unknown", "initializing", "off", "on"] Introduced in: 9.8 |
| supported | boolean | This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field. |
| unsupported_reason | unsupported_reason | |

_links

| Name | Type | Description |
|----------|----------------------|-------------|
| suspects | href | |

anti_ransomware_attack_report

| Name | Type | Description |
|--------|------------------------|---|
| _links | _links | |
| time | string | Timestamp at which ransomware attack is observed. |

space

| Name | Type | Description |
|-------------------|---------|--|
| snapshot_count | integer | Total number of Anti-ransomware backup Snapshot copies. |
| used | integer | Total space in bytes used by the Anti-ransomware feature. |
| used_by_logs | integer | Space in bytes used by the Anti-ransomware analytics logs. |
| used_by_snapshots | integer | Space in bytes used by the Anti-ransomware backup Snapshot copies. |

suspect_files

| Name | Type | Description |
|--------|---------|---|
| count | integer | Total number of <code>suspect_files.format</code> files observed by the Anti-ransomware analytics engine on the volume. |
| format | string | File formats observed by the Anti-ransomware analytics engine on the volume. |

anti_ransomware

Anti-ransomware related information of the volume.

| Name | Type | Description |
|--------------------|--|--|
| attack_probability | string | Probability of a ransomware attack. <code>none</code> No files are suspected of ransomware activity. <code>low</code> A number of files are suspected of ransomware activity. <code>moderate</code> A moderate number of files are suspected of ransomware activity. <code>high</code> A large number of files are suspected of ransomware activity. |
| attack_reports | array[anti_ransomware_attack_report] | |

| Name | Type | Description |
|--------------------|--|--|
| dry_run_start_time | string | Time when Anti-ransomware monitoring <code>state</code> is set to dry-run value for starting evaluation mode. |
| space | space | |
| state | string | <p>Anti-ransomware state.</p> <p><code>disabled</code> Anti-ransomware monitoring is disabled on the volume. This is the default state in a POST operation.</p> <p><code>disable_in_progress</code> Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. Valid in GET operation.</p> <p><code>dry_run</code> Anti-ransomware monitoring is enabled in the evaluation mode.</p> <p><code>enabled</code> Anti-ransomware monitoring is active on the volume.</p> <p><code>paused</code> Anti-ransomware monitoring is paused on the volume.</p> <p><code>enable_paused</code> Anti-ransomware monitoring is paused on the volume from its earlier enabled state. Valid in GET operation.</p> <p><code>dry_run_paused</code> Anti-ransomware monitoring is paused on the volume from its earlier <code>dry_run</code> state. Valid in GET operation. For POST, the valid Anti-ransomware states are only <code>disabled</code>, <code>enabled</code> and <code>dry_run</code>, whereas for PATCH, <code>paused</code> is also valid along with the three valid states for POST.</p> |
| suspect_files | array[suspect_files] | |

application

| Name | Type | Description |
|------|--------|--|
| name | string | Name of the application to which the volume belongs. Available only when the volume is part of an application. |

| Name | Type | Description |
|------|--------|--|
| uuid | string | UUID of the application to which the volume belongs. Available only when the volume is part of an application. |

autosize

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

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

snapshot_reference

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

parent_svm

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

parent_volume

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

| Name | Type | Description |
|------|--------|--|
| uuid | string | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6 |

clone

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

consistency_group

Consistency group the volume is part of.

| Name | Type | Description |
|------|--------|--|
| name | string | The name of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |
| uuid | string | The UUID of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |

aggregates

| Name | Type | Description |
|------|--------|--|
| name | string | Name of the aggregate hosting the FlexGroup Constituent. |
| uuid | string | Unique identifier for the aggregate. |

destination_aggregate

Aggregate

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

movement

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

| Name | Type | Description |
|-----------------------|---------------------------------------|---|
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

logical_space

| Name | Type | Description |
|-------------|---------|--|
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |

| Name | Type | Description |
|-------------|---------|---|
| 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 |
|----------------------------------|---------|---|
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| block_storage_inactive_user_data | integer | The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |

| Name | Type | Description |
|----------------------------|-------------------------------|--|
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| snapshot | snapshot | |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| used_percent | integer | The virtual space used (includes volume reserves) before storage efficiency, as a percent. |

constituents

| Name | Type | Description |
|------------|----------------------------|-------------|
| aggregates | aggregates | |

| 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 | FlexGroup Constituents name |
| space | space | |

policy

| Name | Type | Description |
|------|--------|---|
| name | string | Specifies the name of the efficiency policy. Valid for PATCH. |

efficiency

| Name | Type | Description |
|---------------------|--------|--|
| application_io_size | string | Block size to use by compression. Valid for POST. |
| 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. |

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

| 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. |
| last_op_begin | string | Last sis operation begin timestamp. |
| last_op_end | string | Last sis operation end timestamp. |
| last_op_err | string | Last sis operation error text. |
| last_op_size | integer | Last sis operation size. |
| last_op_state | string | Last sis operation state. |
| op_state | string | Sis status of the volume. |
| path | string | Absolute volume path of the volume. |
| policy | policy | |
| progress | string | Sis progress of the volume. |
| schedule | string | Schedule associated with volume. |

| Name | Type | Description |
|-------------------------|--------|--|
| state | string | Storage efficiency state of the volume. Currently, this field supports POST/PATCH only for RW (Read-Write) volumes on FSx for ONTAP. disabled ‐ All storage efficiency features are disabled. mixed ‐ Read-only field for FlexGroup volumes, storage efficiency is enabled on certain constituents and disabled on others. On FSx for ONTAP ‐ enabled ‐ All supported storage efficiency features for the volume are enabled. custom ‐ Read-only field currently only supported for the FSx for ONTAP, user-defined storage efficiency features are enabled. For other platforms ‐ enabled ‐ At least one storage efficiency feature for the volume is enabled. <ul style="list-style-type: none"> enum: ["disabled", "enabled", "mixed", "custom"] Introduced in: 9.9 |
| storage_efficiency_mode | string | Storage efficiency mode used by volume. This parameter is supported only on AFF platform. |
| type | string | Sis Type of the volume. |

status

| Name | Type | Description |
|---------|--------|-----------------------------------|
| code | string | Encryption progress message code. |
| message | string | Encryption progress message. |

encryption

| Name | Type | Description |
|-----------------------|---------|--|
| enabled | boolean | Creates an encrypted or an unencrypted volume. For POST, when set to 'true', a new key is generated and used to encrypt the given volume. In that case, the underlying SVM must be configured with the key manager. When set to 'false', the volume created will be unencrypted. For PATCH, when set to 'true', it encrypts an unencrypted volume. Specifying the parameter as 'false' in a PATCH operation for an encrypted volume is only supported when moving the volume to another aggregate. |
| key_id | string | The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key. |
| key_manager_attribute | string | Specifies an additional key manager attribute that is an identifier-value pair, separated by '='. For example, CRN=unique-value. This parameter is required when using the POST method and an IBM Key Lore key manager is configured on the SVM. |
| rekey | boolean | If set to 'true', re-encrypts the volume with a new key. Valid in PATCH. |
| state | string | Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one. |

| Name | Type | Description |
|--------|------------------------|---|
| status | status | |
| type | string | Volume encryption type. none ‐ The volume is a plain-text one. volume ‐ The volume is encrypted with NVE (NetApp Volume Encryption). aggregate ‐ The volume is encrypted with NAE (NetApp Aggregate Encryption). |

error_state

| Name | Type | Description |
|-----------------|---------|---|
| has_bad_blocks | boolean | Indicates whether the volume has any corrupt data blocks. If the damaged data block is accessed, an IO error, such as EIO for NFS or STATUS_FILE_CORRUPT for CIFS, is returned. |
| is_inconsistent | boolean | Indicates whether the file system has any inconsistencies. true ‐ File system is inconsistent. false ‐ File system 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. |

flash_pool

| Name | Type | Description |
|--------------------------|--------|---|
| cache_eligibility | string | If this parameter is specified, the command displays information only about the volume or volumes with the specified Flash Pool caching attributes. |
| cache_retention_priority | string | If this parameter is specified, the command displays the volumes that match the specified cache retention priority policy. A cache retention priority defines how long the blocks of a volume will be cached in the Flash Pool once they become cold. |

| Name | Type | Description |
|----------------|--------|---|
| caching_policy | string | This optionally specifies the caching policy to apply to the volume. A caching policy defines how the system caches a volume's data in Flash Cache modules. If a caching policy is not assigned to a volume, the system uses the caching policy that is assigned to the containing SVM. If a caching policy is not assigned to the containing SVM, the system uses the default cluster-wide policy. |

flexgroup

| Name | Type | Description |
|------|--------|---|
| name | string | Name of the FlexGroup volume that the constituent is part of. |
| uuid | string | Unique identifier for the FlexGroup volume that the constituent is part of. |

guarantee

| Name | Type | Description |
|---------|---------|---|
| honored | boolean | Is the space guarantee of this volume honored in the aggregate? |
| type | string | The type of space guarantee of this volume in the aggregate. |

iops

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

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

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

latency

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

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

cloud

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.

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

| Name | Type | Description |
|-----------|---------|---|
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

flexcache

Performance number for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
|--------------------|---------|--|
| bandwidth_savings | integer | Bandwidth savings denoting the amount of data served locally by the cache, in bytes. |
| cache_miss_percent | integer | Cache miss percentage. |

| Name | Type | Description |
|-----------|--------|---|
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

throughput

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

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

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

metric

Performance numbers, such as IOPS, latency and throughput.

| Name | Type | Description |
|-----------|---------------------------|--|
| _links | _links | |
| cloud | cloud | Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| flexcache | flexcache | Performance number for FlexCache used to measure cache effectiveness. |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
|------------|----------------------------|---|
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

movement

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

| Name | Type | Description |
|-----------------------|---------------------------------------|---|
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| start_time | string | Start time of volume move. |

| Name | Type | Description |
|----------------|--------|---|
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

export_policy

Export Policy

| Name | Type | Description |
|--------|------------------------|-------------|
| _links | _links | |
| id | integer | |
| name | string | |

junction_parent

| Name | Type | Description |
|--------|------------------------|--|
| _links | _links | |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |
| uuid | string | Unique identifier for the parent volume. |

nas

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

| Name | Type | Description |
|------------------|---------|---|
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth digit), first digit is assumed to be zero. |

policy

| Name | Type | Description |
|------------------------|------------------------|---|
| _links | _links | |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. |

| Name | Type | Description |
|---------------------|---------|---|
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

QoS information

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

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

| Name | Type | Description |
|---------|---------|--|
| enabled | boolean | This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVols or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVols or FlexGroup volumes when the quota state is "off". |
| state | string | Quota state of the volume |

retention

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

destinations

| Name | Type | Description |
|----------|---------|---|
| is_cloud | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to a cloud destination. |
| is_ontap | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to an ONTAP destination. <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.9 |

snapmirror

Specifies attributes for SnapMirror protection.

| Name | Type | Description |
|--------------|------------------------------|---|
| destinations | destinations | |
| is_protected | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data. |

snapshot_policy

This is a reference to the Snapshot copy policy.

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

logical_space

| Name | Type | Description |
|-------------|---------|---|
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |
| used | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), in bytes. |
| used_by_afs | integer | The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes. |

| Name | Type | Description |
|-------------------|---------|--|
| used_by_snapshots | integer | Size that is logically used across all Snapshot copies in the volume, in bytes. |
| used_percent | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), as a percentage. |

snapshot

| Name | Type | Description |
|--------------------|---------|---|
| autodelete_enabled | boolean | Specifies whether Snapshot copy autodelete is currently enabled on this volume. |
| autodelete_trigger | string | Specifies when the system should trigger an autodelete of Snapshot copies. When set to <i>volume</i> , autodelete is triggered based on volume fullness. When set to <i>snap_reserve</i> , autodelete is triggered based on Snapshot reserve fullness. The default value is <i>volume</i> . |
| reserve_available | integer | Size available for Snapshot copies within the Snapshot copy reserve, in bytes. |
| reserve_percent | integer | The space that has been set aside as a reserve for Snapshot copy usage, in percent. |
| reserve_size | integer | Size in the volume that has been set aside as a reserve for Snapshot copy usage, in bytes. |
| space_used_percent | integer | Percentage of snapshot reserve size that has been used. |
| used | integer | The total space used by Snapshot copies in the volume, in bytes. |

space

| Name | Type | Description |
|---|---------|---|
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| 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. |
| block_storage_inactive_user_data_percent | integer | Percentage of size that is physically used in the performance tier of the volume. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |
| cross_volume_dedupe_metafiles_footprint | integer | Cross volume deduplication metadata footprint, in bytes. |
| cross_volume_dedupe_metafiles_temporary_footprint | integer | Cross volume temporary deduplication metadata footprint, in bytes. |
| dedupe_metafiles_footprint | integer | Deduplication metadata footprint, in bytes. |
| dedupe_metafiles_temporary_footprint | integer | Temporary deduplication metadata footprint, in bytes. |
| delayed_free_footprint | integer | Delayed free blocks footprint, in bytes. |
| expected_available | integer | Size that should be available for the volume, irrespective of available size in the aggregate, in bytes. |

| Name | Type | Description |
|-------------------------------|-------------------------------|---|
| file_operation_metadata | integer | File operation metadata footprint, in bytes. |
| filesystem_size | integer | Total usable size of the volume, in bytes. |
| filesystem_size_fixed | boolean | Specifies whether the file system is to remain of the same size when set to true or to grow when set to false. This option is automatically set to true when a volume becomes SnapMirrored. |
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| fractional_reserve | integer | Used to change the amount of space reserved for overwrites of reserved objects in a volume. |
| full_threshold_percent | integer | Volume full threshold percentage at which EMS warnings can be sent. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| nearly_full_threshold_percent | integer | Volume nearly full threshold percentage at which EMS warnings can be sent. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| overwrite_reserve | integer | Reserved space for overwrites, in bytes. |
| overwrite_reserve_used | integer | Overwrite logical reserve space used, in bytes. |

| Name | Type | Description |
|----------------------------------|--------------------------|---|
| percent_used | integer | Percentage of the volume size that is used. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| physical_used | integer | Size that is physically used in the volume, in bytes. |
| physical_used_percent | integer | Size that is physically used in the volume, as a percentage. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| size_available_for_snapshots | integer | Available space for Snapshot copies from snap-reserve, in bytes. |
| snapmirror_destination_footprint | integer | SnapMirror destination footprint, in bytes. |
| snapshot | snapshot | |
| snapshot_reserve_unusable | integer | Snapshot reserve that is not available for Snapshot copy creation, in bytes. |
| snapshot_spill | integer | Space used by the snapshot copies beyond the snap-reserve, in bytes. |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| user_data | integer | User data, in bytes. |

| Name | Type | Description |
|----------------------------|---------|---|
| volume_guarantee_footprint | integer | Space reserved for future writes in the volume, in bytes. |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

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

latency_raw

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

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

cloud

These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.

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

flexcache_raw

Performance numbers for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
|-------------------------|---------|---|
| cache_miss_blocks | integer | Blocks retrieved from origin in case of a cache miss. This can be divided by the raw client_requested_blocks and multiplied by 100 to calculate the cache miss percentage. |
| client_requested_blocks | integer | Total blocks requested by the client. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

throughput_raw

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

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

statistics

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

| Name | Type | Description |
|---------------|-------------------------------|--|
| cloud | cloud | These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools. |
| flexcache_raw | flexcache_raw | Performance numbers for FlexCache used to measure cache effectiveness. |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |

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

svm

SVM containing the volume. Required on POST.

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

tiering

| Name | Type | Description |
|------------------|---------------|--|
| min_cooling_days | integer | <p>This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. Note that this parameter is only used for tiering purposes and does not affect the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data. This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy.</p> |
| object_tags | array[string] | <p>This parameter specifies tags of a volume for objects stored on a FabricPool-enabled aggregate. Each tag is a key,value pair and should be in the format "key=value".</p> |

| Name | Type | Description |
|--------|--------|--|
| policy | string | <p>Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.</p> <p>all &dash; This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.</p> <p>auto &dash; This policy allows tiering of both snapshot and active file system user data to the cloud store</p> <p>none &dash; Volume blocks will not be tiered to the cloud store.</p> <p>snapshot_only &dash; This policy allows tiering of only the volume Snapshot copies not associated with the active file system.</p> <p>The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days.</p> |

| Name | Type | Description |
|-----------|---------|---|
| supported | boolean | This parameter specifies whether or not FabricPools are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only FabricPool aggregates are used if this parameter is set to true and only non FabricPool aggregates are used if this parameter is set to false. Tiering support for a FlexGroup can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only does the volume need to support tiering by using FabricPools, the tiering "policy" must not be 'none'. A volume that uses FabricPools but has a tiering "policy" of 'none' supports tiering, but will not tier any data. |

volume

| Name | Type | Description |
|---------------------|-------------------------------------|---|
| _links | _links | |
| access_time_enabled | boolean | Indicates whether or not access time updates are enabled on the volume. |
| activity_tracking | activity_tracking | |
| aggregates | array[aggregates] | Aggregate hosting the volume. Required on POST. |
| analytics | analytics | |
| anti_ransomware | anti_ransomware | Anti-ransomware related information of the volume. |

| Name | Type | Description |
|-----------------------|-----------------------------|---|
| anti_ransomware_state | string | <p>The Anti-ransomware state of the volume. If no "anti_ransomware_state" property is specified, the volume inherits the value from its parent SVM's "anti_ransomware_default_volume_state" property. If this value is "disabled", Anti-ransomware is disabled on the volume. If this value is "enabled", Anti-ransomware is enabled on the volume and alerts are raised if any suspect is detected for those volumes. If this value is "dry_run", Anti-ransomware is enabled in the dry-run or learning mode on the volume. The "dry_run" state is same as the "enabled" state except that the analytics data is used here for learning. No alerts are raised for any detections or violations. If this value is "paused", Anti-ransomware is paused on the volume. Additionally, three more states are available, which are only valid for GET. If this value is "disable_in_progress", Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. If this value is "enable_paused", Anti-ransomware is paused on the volume from its earlier enabled state. If this value is "dry_run_paused", Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. For POST, the valid Anti-ransomware states are only "disabled", "enabled" and "dry_run", whereas for PATCH, "paused" is also valid along with the three valid states for POST.</p> |
| application | application | |
| autosize | autosize | |
| clone | clone | |

| Name | Type | Description |
|------------------------|---------------------------------------|---|
| cloud_retrieval_policy | string | <p>This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved. "on_read" policy retrieves tiered data for all client driven data reads. "never" policy never retrieves tiered data. "promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved.</p> |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| consistency_group | consistency_group | Consistency group the volume is part of. |
| constituents | array[constituents] | |

| 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 volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list. |
| convert_unicode | boolean | Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients. |
| 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 | |
| flash_pool | flash_pool | |
| 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. |
| flexgroup | flexgroup | |
| guarantee | guarantee | |

| Name | Type | Description |
|-----------------|--------------------------|--|
| is_object_store | boolean | Specifies whether the volume is provisioned for an object store server. |
| is_svm_root | boolean | Specifies whether the volume is a root volume of the SVM it belongs to. |
| language | string | Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting. |
| max_dir_size | integer | Maximum directory size. This value sets maximum size, in bytes, to which a directory can grow. The default maximum directory size for FlexVol volumes is model-dependent, and optimized for the size of system memory. Before increasing the maximum directory size, involve technical support. |
| 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 | |

| Name | Type | Description |
|----------------------------------|-----------------|---|
| qos | qos | QoS information |
| queue_for_encryption | boolean | Specifies whether the volume is queued for encryption. |
| quota | quota | Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume. |
| scheduled_snapshot_naming_scheme | string | <p>Naming Scheme for automatic Snapshot copies:</p> <ul style="list-style-type: none"> • create_time - Automatic Snapshot copies are saved as per the start of their current date and time. • ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> |
| size | integer | Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size. |
| snaplock | snaplock | |
| snapmirror | snapmirror | Specifies attributes for SnapMirror protection. |
| snapshot_count | integer | Number of Snapshot copies in the volume. |
| snapshot_policy | snapshot_policy | This is a reference to the Snapshot copy policy. |
| space | space | |

| Name | Type | Description |
|------------|----------------------------|---|
| state | string | Volume state. Client access is supported only when volume is online and junctioned. Taking volume to offline or restricted state removes its junction path and blocks client access. When volume is in restricted state some operations like parity reconstruction and iron on commit are allowed. 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. |
| status | array[string] | Describes the current status of a volume. |

| Name | Type | Description |
|---------|---------|---|
| style | string | <p>The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol" you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. The style "flexgroup_constituent" is not supported when creating a volume. flexvol &dash; flexible volumes and FlexClone volumes flexgroup &dash; FlexGroup volumes flexgroup_constituent &dash; FlexGroup constituents.</p> |
| svm | svm | SVM containing the volume. Required on POST. |
| tiering | tiering | |
| type | string | <p>Type of the volume. rw &dash; read-write volume. dp &dash; data-protection volume. ls &dash; load-sharing <code>dp</code> volume. Valid in GET.</p> |

| Name | Type | Description |
|-------------------------|---------|--|
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration. |
| uuid | string | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6 |

job_link

| Name | Type | Description |
|------------------------|------------------------|---|
| _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}

Introduced In: 9.6

Deletes a volume. If the UUID belongs to a volume, all of its blocks are freed and returned to its containing aggregate. If a volume is online, it is offlined before deletion. If a volume is mounted, unmount the volume by specifying the nas.path as empty before deleting it using the DELETE operation.

Related ONTAP commands

- `volume delete`
- `volume clone delete`

Parameters

| Name | Type | In | Required | Description |
|------|--------|------|----------|----------------------------------|
| uuid | string | path | True | Unique identifier of the volume. |

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

Response

Status: 202, Accepted

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

Example response

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

Error

Status: Default, Error

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

Example error

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

Definitions

See Definitions

href

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

_links

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

job_link

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

error_arguments

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

error

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

Retrieve a volume

GET /storage/volumes/{uuid}

Introduced In: 9.6

Retrieves a volume. The GET API can be used to retrieve the quota state for a FlexVol or a FlexGroup volume.

Expensive properties

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

- `is_svm_root`
- `analytics.*`
- `anti_ransomware.*`
- `application.*`
- `encryption.*`
- `convert_unicode`
- `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.*`
- `max_dir_size`
- `nas.export_policy.id`
- `nas.gid`
- `nas.path`
- `nas.security_style`
- `nas.uid`
- `nas.unix_permissions`
- `nas.junction_parent.name`
- `nas.junction_parent.uuid`

- snaplock.*
- restore_to.*
- snapshot_policy.uuid
- quota.*
- qos.*
- flexcache_endpoint_type
- space.block_storage_inactive_user_data
- space.capacity_tier_footprint
- space.performance_tier_footprint
- space.local_tier_footprint
- space.footprint
- space.over_provisioned
- space.metadata
- space.total_footprint
- space.dedupe_metafiles_footprint
- space.dedupe_metafiles_temporary_footprint
- space.delayed_free_footprint
- space.file_operation_metadata
- space.snapmirror_destination_footprint
- space.volume_guarantee_footprint
- space.cross_volume_dedupe_metafiles_footprint
- space.cross_volume_dedupe_metafiles_temporary_footprint
- space.snapshot_reserve_unusable
- space.snapshot_spill
- space.user_data
- space.logical_space.*
- space.snapshot.*
- space.used_by_afs
- space.afs_total
- space.available_percent
- space.full_threshold_percent
- space.nearly_full_threshold_percent
- space.overwrite_reserve
- space.overwrite_reserve_used

- `space.size_available_for_snapshots`
- `space.percent_used`
- `space.fractional_reserve`
- `space.block_storage_inactive_user_data_percent`
- `space.physical_used`
- `space.physical_used_percent`
- `space.expected_available`
- `space.filesystem_size`
- `space.filesystem_size_fixed`
- `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`
- `security anti-ransomware volume show`
- `security anti-ransomware volume attack generate-report`
- `security anti-ransomware volume space show`

Parameters

| Name | Type | In | Required | Description |
|-------------------|---------------------|-------------------|----------|----------------------------------|
| <code>uuid</code> | <code>string</code> | <code>path</code> | True | Unique identifier of the volume. |

| Name | Type | In | Required | Description |
|----------------|---------------|-------|----------|--|
| is_constituent | boolean | query | False | When set to false, only FlexVol and FlexGroup volumes are returned. When set to true, only FlexGroup constituent volumes are returned. Default for GET calls is false. <ul style="list-style-type: none"> • Introduced in: 9.10 • Default value: |
| fields | array[string] | query | False | Specify the fields to return. |

Response

Status: 200, Ok

| Name | Type | Description |
|---------------------|-------------------------------------|---|
| _links | _links | |
| access_time_enabled | boolean | Indicates whether or not access time updates are enabled on the volume. |
| activity_tracking | activity_tracking | |
| aggregates | array[aggregates] | Aggregate hosting the volume. Required on POST. |
| analytics | analytics | |
| anti_ransomware | anti_ransomware | Anti-ransomware related information of the volume. |

| Name | Type | Description |
|-----------------------|-------------|---|
| anti_ransomware_state | string | <p>The Anti-ransomware state of the volume. If no "anti_ransomware_state" property is specified, the volume inherits the value from its parent SVM's "anti_ransomware_default_volume_state" property. If this value is "disabled", Anti-ransomware is disabled on the volume. If this value is "enabled", Anti-ransomware is enabled on the volume and alerts are raised if any suspect is detected for those volumes. If this value is "dry_run", Anti-ransomware is enabled in the dry-run or learning mode on the volume. The "dry_run" state is same as the "enabled" state except that the analytics data is used here for learning. No alerts are raised for any detections or violations. If this value is "paused", Anti-ransomware is paused on the volume. Additionally, three more states are available, which are only valid for GET. If this value is "disable_in_progress", Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. If this value is "enable_paused", Anti-ransomware is paused on the volume from its earlier enabled state. If this value is "dry_run_paused", Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. For POST, the valid Anti-ransomware states are only "disabled", "enabled" and "dry_run", whereas for PATCH, "paused" is also valid along with the three valid states for POST.</p> |
| application | application | |
| autosize | autosize | |
| clone | clone | |

| Name | Type | Description |
|------------------------|---------------------------------------|---|
| cloud_retrieval_policy | string | <p>This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved. "on_read" policy retrieves tiered data for all client driven data reads. "never" policy never retrieves tiered data. "promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved.</p> |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| consistency_group | consistency_group | Consistency group the volume is part of. |
| constituents | array[constituents] | |

| 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 volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list. |
| convert_unicode | boolean | Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients. |
| 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 | |
| flash_pool | flash_pool | |
| 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. |
| flexgroup | flexgroup | |
| guarantee | guarantee | |
| is_object_store | boolean | Specifies whether the volume is provisioned for an object store server. |

| Name | Type | Description |
|----------------------|--------------------------|--|
| is_svm_root | boolean | Specifies whether the volume is a root volume of the SVM it belongs to. |
| language | string | Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting. |
| max_dir_size | integer | Maximum directory size. This value sets maximum size, in bytes, to which a directory can grow. The default maximum directory size for FlexVol volumes is model-dependent, and optimized for the size of system memory. Before increasing the maximum directory size, involve technical support. |
| 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 |
| queue_for_encryption | boolean | Specifies whether the volume is queued for encryption. |

| Name | Type | Description |
|----------------------------------|---------------------------------|---|
| quota | quota | Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume. |
| scheduled_snapshot_naming_scheme | string | <p>Naming Scheme for automatic Snapshot copies:</p> <ul style="list-style-type: none"> • create_time - Automatic Snapshot copies are saved as per the start of their current date and time. • ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> |
| size | integer | Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size. |
| snaplock | snaplock | |
| snapmirror | snapmirror | Specifies attributes for SnapMirror protection. |
| snapshot_count | integer | Number of Snapshot copies in the volume. |
| snapshot_policy | snapshot_policy | This is a reference to the Snapshot copy policy. |
| space | space | |

| Name | Type | Description |
|------------|----------------------------|---|
| state | string | Volume state. Client access is supported only when volume is online and junctioned. Taking volume to offline or restricted state removes its junction path and blocks client access. When volume is in restricted state some operations like parity reconstruction and iron on commit are allowed. 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. |
| status | array[string] | Describes the current status of a volume. |

| Name | Type | Description |
|---------|---------|--|
| style | string | <p>The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol" you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. The style "flexgroup_constituent" is not supported when creating a volume.</p> <p>flexvol &dash; flexible volumes and FlexClone volumes flexgroup &dash; FlexGroup volumes flexgroup_constituent &dash; FlexGroup constituents.</p> |
| svm | svm | SVM containing the volume. Required on POST. |
| tiering | tiering | |
| type | string | <p>Type of the volume. rw &dash; read-write volume. dp &dash; data-protection volume. ls &dash; load-sharing <code>dp</code> volume. Valid in GET.</p> |

| 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 • Introduced in: 9.6 |

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "activity_tracking": {
    "state": "off",
    "unsupported_reason": {
      "code": "124518405",
      "message": "Volume activity tracking cannot be enabled on volumes
that contain LUNs."
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "analytics": {
    "scan_progress": 17,
    "state": "unknown",
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
    }
  },
  "anti_ransomware": {
    "attack_probability": "none",
    "attack_reports": {
      "_links": {
        "suspects": {
          "href": "/api/resourcelink"
        }
      }
    },
    "time": "2021-06-01T20:36:41+05:30"
  },
  "dry_run_start_time": "string",
  "space": {
```

```

    "snapshot_count": 0,
    "used": 0,
    "used_by_logs": 0,
    "used_by_snapshots": 0
  },
  "state": "disabled",
  "suspect_files": {
    "count": 0,
    "format": "string"
  }
},
"anti_ransomware_state": "disabled",
"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
  },
  "cloud_retrieval_policy": "default",
  "comment": "string",
  "consistency_group": {
    "name": "consistency_group_1",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "constituents": {
    "aggregates": {
      "name": "string",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "movement": {
      "cutover_window": 30,
      "destination_aggregate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "percent_complete": 0,
      "state": "replicating",
      "tiering_policy": "all"
    },
    "name": "string",
    "space": {
      "available": 0,
      "block_storage_inactive_user_data": 0,
      "capacity_tier_footprint": 0,
      "footprint": 0,
      "local_tier_footprint": 0,
      "logical_space": {
        "available": 0,
        "used_by_afs": 0
      },
      "metadata": 0,
      "over_provisioned": 0,
      "performance_tier_footprint": 0,
      "snapshot": {
        "used": 0
      }
    }
  }
}

```

```

    },
    "total_footprint": 0,
    "used": 0
  }
},
"create_time": "2018-06-04T19:00:00Z",
"efficiency": {
  "application_io_size": "8k",
  "compaction": "inline",
  "compression": "inline",
  "cross_volume_dedupe": "inline",
  "dedupe": "inline",
  "last_op_begin": "string",
  "last_op_end": "string",
  "last_op_err": "string",
  "last_op_size": 0,
  "last_op_state": "string",
  "op_state": "idle",
  "path": "string",
  "progress": "string",
  "schedule": "string",
  "state": "disabled",
  "storage_efficiency_mode": "default",
  "type": "regular"
},
"encryption": {
  "key_id": "string",
  "key_manager_attribute": "CRN=v1:bluemix:public:containers-
kubernetes:us-south:a/asdfghjkl1234:asdfghjkl1234:worker:kubernetes-
asdfghjkl-worker1",
  "state": "encrypted",
  "status": {
    "code": "string",
    "message": "string"
  },
  "type": "none"
},
"files": {
  "used": 0
},
"flash_pool": {
  "cache_eligibility": "read",
  "cache_retention_priority": "normal",
  "caching_policy": "none"
},
"flexcache_endpoint_type": "none",

```

```
"flexgroup": {
  "name": "my_flexgroup",
  "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
},
"guarantee": {
  "type": "volume"
},
"language": "ar",
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"cloud": {
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "timestamp": "2017-01-25T11:20:13Z"
},
"duration": "PT15S",
"flexcache": {
  "bandwidth_savings": 4096,
  "cache_miss_percent": 20,
  "duration": "PT1D",
  "status": "ok",
  "timestamp": "2017-01-25T11:20:13Z"
},
"iops": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"latency": {
  "read": 200,
  "total": 1000,
  "write": 100
}
```

```

    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": 0,
    "start_time": "2020-12-07T03:45:12-05:00",
    "state": "replicating",
    "tiering_policy": "all"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 100,
      "name": "default"
    },
    "junction_parent": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "vs1_root",
      "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
    },
    "path": "/user/my_volume",

```

```

    "security_style": "mixed",
    "unix_permissions": 755
  },
  "qos": {
    "policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    },
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"quota": {
  "state": "corrupt"
},
"scheduled_snapshot_naming_scheme": "create_time",
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04T19:00:00Z",
  "expiry_time": "Wed Sep 5 11:02:42 GMT 2018",
  "is_audit_log": 1,
  "litigation_count": 10,
  "privileged_delete": "enabled",
  "retention": {
    "default": "P30Y",
    "maximum": "P30Y",
    "minimum": "P30Y"
  },
  "type": "enterprise",
  "unspecified_retention_file_count": 10
},
"snapshot_count": 0,
"snapshot_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"name": "default",

```

```
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"space": {
  "available": 0,
  "block_storage_inactive_user_data": 0,
  "block_storage_inactive_user_data_percent": 0,
  "capacity_tier_footprint": 0,
  "cross_volume_dedupe_metafiles_footprint": 0,
  "cross_volume_dedupe_metafiles_temporary_footprint": 0,
  "dedupe_metafiles_footprint": 0,
  "dedupe_metafiles_temporary_footprint": 0,
  "delayed_free_footprint": 0,
  "file_operation_metadata": 0,
  "filesystem_size": 0,
  "footprint": 0,
  "local_tier_footprint": 0,
  "logical_space": {
    "available": 0,
    "used": 0,
    "used_by_afs": 0,
    "used_by_snapshots": 0,
    "used_percent": 0
  },
  "metadata": 0,
  "over_provisioned": 0,
  "overwrite_reserve": 0,
  "overwrite_reserve_used": 0,
  "percent_used": 0,
  "performance_tier_footprint": 0,
  "size_available_for_snapshots": 0,
  "snapmirror_destination_footprint": 0,
  "snapshot": {
    "autodelete_trigger": "volume",
    "reserve_available": 0,
    "reserve_size": 0,
    "space_used_percent": 0,
    "used": 0
  },
  "snapshot_reserve_unusable": 0,
  "snapshot_spill": 0,
  "total_footprint": 0,
  "used": 0,
  "user_data": 0,
  "volume_guarantee_footprint": 0
},
"state": "error",
```

```
"statistics": {
  "cloud": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "flexcache_raw": {
    "cache_miss_blocks": 10,
    "client_requested_blocks": 500,
    "status": "ok",
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "iops_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"status": {
},
"style": "flexvol",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
```

```

    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"tiering": {
  "object_tags": {
  },
  "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 | |

unsupported_reason

| Name | Type | Description |
|---------|--------|--|
| code | string | If volume activity tracking is not supported on the volume, this field provides an appropriate error code. |
| message | string | If volume activity tracking is not supported on the volume, this field provides an error message detailing why this is the case. |

activity_tracking

| Name | Type | Description |
|-------|--------|---|
| state | string | Activity tracking state of the volume. If this value is "on", ONTAP collects top metrics information for the volume in real time. There is a slight impact to I/O performance in order to collect this information. If this value is "off", no activity tracking information is collected or available to view. <ul style="list-style-type: none">• enum: ["off", "on"]• Introduced in: 9.10 |

| Name | Type | Description |
|--------------------|------------------------------------|--|
| supported | boolean | This field indicates whether or not volume activity tracking is supported on the volume. If volume activity tracking is not supported, the reason why is provided in the "activity_tracking.unsupported_reason" field. |
| unsupported_reason | unsupported_reason | |

aggregates

Aggregate

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

unsupported_reason

| Name | Type | Description |
|---------|--------|--|
| code | string | If file system analytics is not supported on the volume, this field provides the error code explaining why. |
| message | string | If file system analytics is not supported on the volume, this field provides the error message explaining why. |

analytics

| Name | Type | Description |
|---------------|---------|---|
| scan_progress | integer | Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is <code>initializing</code> . |

| Name | Type | Description |
|--------------------|------------------------------------|---|
| state | string | File system analytics state of the volume. If this value is "on", ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is "off", file system analytics information is not collected and not available to be viewed. If this value is "initializing", that means file system analytics was recently turned on, and the initialization scan to gather information all existing files and directories is currently running. If this value is 'unknown' that means there was an internal error when determining the file system analytics state for the volume. <ul style="list-style-type: none"> enum: ["unknown", "initializing", "off", "on"] Introduced in: 9.8 |
| supported | boolean | This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field. |
| unsupported_reason | unsupported_reason | |

_links

| Name | Type | Description |
|----------|----------------------|-------------|
| suspects | href | |

anti_ransomware_attack_report

| Name | Type | Description |
|--------|------------------------|---|
| _links | _links | |
| time | string | Timestamp at which ransomware attack is observed. |

space

| Name | Type | Description |
|-------------------|---------|--|
| snapshot_count | integer | Total number of Anti-ransomware backup Snapshot copies. |
| used | integer | Total space in bytes used by the Anti-ransomware feature. |
| used_by_logs | integer | Space in bytes used by the Anti-ransomware analytics logs. |
| used_by_snapshots | integer | Space in bytes used by the Anti-ransomware backup Snapshot copies. |

suspect_files

| Name | Type | Description |
|--------|---------|---|
| count | integer | Total number of <code>suspect_files.format</code> files observed by the Anti-ransomware analytics engine on the volume. |
| format | string | File formats observed by the Anti-ransomware analytics engine on the volume. |

anti_ransomware

Anti-ransomware related information of the volume.

| Name | Type | Description |
|--------------------|--|--|
| attack_probability | string | Probability of a ransomware attack. <code>none</code> No files are suspected of ransomware activity. <code>low</code> A number of files are suspected of ransomware activity. <code>moderate</code> A moderate number of files are suspected of ransomware activity. <code>high</code> A large number of files are suspected of ransomware activity. |
| attack_reports | array[anti_ransomware_attack_report] | |

| Name | Type | Description |
|--------------------|--|--|
| dry_run_start_time | string | Time when Anti-ransomware monitoring <code>state</code> is set to dry-run value for starting evaluation mode. |
| space | space | |
| state | string | <p>Anti-ransomware state.</p> <p><code>disabled</code> Anti-ransomware monitoring is disabled on the volume. This is the default state in a POST operation.</p> <p><code>disable_in_progress</code> Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. Valid in GET operation.</p> <p><code>dry_run</code> Anti-ransomware monitoring is enabled in the evaluation mode.</p> <p><code>enabled</code> Anti-ransomware monitoring is active on the volume.</p> <p><code>paused</code> Anti-ransomware monitoring is paused on the volume.</p> <p><code>enable_paused</code> Anti-ransomware monitoring is paused on the volume from its earlier enabled state. Valid in GET operation.</p> <p><code>dry_run_paused</code> Anti-ransomware monitoring is paused on the volume from its earlier <code>dry_run</code> state. Valid in GET operation. For POST, the valid Anti-ransomware states are only <code>disabled</code>, <code>enabled</code> and <code>dry_run</code>, whereas for PATCH, <code>paused</code> is also valid along with the three valid states for POST.</p> |
| suspect_files | array[suspect_files] | |

application

| Name | Type | Description |
|------|--------|--|
| name | string | Name of the application to which the volume belongs. Available only when the volume is part of an application. |

| Name | Type | Description |
|------|--------|--|
| uuid | string | UUID of the application to which the volume belongs. Available only when the volume is part of an application. |

autosize

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

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

snapshot_reference

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

parent_svm

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

parent_volume

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

| Name | Type | Description |
|------|--------|--|
| uuid | string | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6 |

clone

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

consistency_group

Consistency group the volume is part of.

| Name | Type | Description |
|------|--------|--|
| name | string | The name of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |
| uuid | string | The UUID of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |

aggregates

| Name | Type | Description |
|------|--------|--|
| name | string | Name of the aggregate hosting the FlexGroup Constituent. |
| uuid | string | Unique identifier for the aggregate. |

destination_aggregate

Aggregate

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

movement

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

| Name | Type | Description |
|-----------------------|---------------------------------------|---|
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

logical_space

| Name | Type | Description |
|-------------|---------|--|
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |

| Name | Type | Description |
|-------------|---------|---|
| 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 |
|----------------------------------|---------|---|
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| block_storage_inactive_user_data | integer | The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |

| Name | Type | Description |
|----------------------------|-------------------------------|--|
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| snapshot | snapshot | |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| used_percent | integer | The virtual space used (includes volume reserves) before storage efficiency, as a percent. |

constituents

| Name | Type | Description |
|------------|----------------------------|-------------|
| aggregates | aggregates | |

| 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 | FlexGroup Constituents name |
| space | space | |

policy

| Name | Type | Description |
|------|--------|---|
| name | string | Specifies the name of the efficiency policy. Valid for PATCH. |

efficiency

| Name | Type | Description |
|---------------------|--------|--|
| application_io_size | string | Block size to use by compression. Valid for POST. |
| 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. |

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

| 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. |
| last_op_begin | string | Last sis operation begin timestamp. |
| last_op_end | string | Last sis operation end timestamp. |
| last_op_err | string | Last sis operation error text. |
| last_op_size | integer | Last sis operation size. |
| last_op_state | string | Last sis operation state. |
| op_state | string | Sis status of the volume. |
| path | string | Absolute volume path of the volume. |
| policy | policy | |
| progress | string | Sis progress of the volume. |
| schedule | string | Schedule associated with volume. |

| Name | Type | Description |
|-------------------------|--------|---|
| state | string | <p>Storage efficiency state of the volume. Currently, this field supports POST/PATCH only for RW (Read-Write) volumes on FSx for ONTAP. disabled &dash; All storage efficiency features are disabled. mixed &dash; Read-only field for FlexGroup volumes, storage efficiency is enabled on certain constituents and disabled on others. On FSx for ONTAP &dash; &nbsp; enabled &dash; All supported storage efficiency features for the volume are enabled. &nbsp; custom &dash; Read-only field currently only supported for the FSx for ONTAP, user-defined storage efficiency features are enabled. For other platforms &dash; &nbsp; enabled &dash; At least one storage efficiency feature for the volume is enabled.</p> <ul style="list-style-type: none"> • enum: ["disabled", "enabled", "mixed", "custom"] • Introduced in: 9.9 |
| storage_efficiency_mode | string | Storage efficiency mode used by volume. This parameter is supported only on AFF platform. |
| type | string | Sis Type of the volume. |

status

| Name | Type | Description |
|---------|--------|-----------------------------------|
| code | string | Encryption progress message code. |
| message | string | Encryption progress message. |

encryption

| Name | Type | Description |
|-----------------------|---------|--|
| enabled | boolean | Creates an encrypted or an unencrypted volume. For POST, when set to 'true', a new key is generated and used to encrypt the given volume. In that case, the underlying SVM must be configured with the key manager. When set to 'false', the volume created will be unencrypted. For PATCH, when set to 'true', it encrypts an unencrypted volume. Specifying the parameter as 'false' in a PATCH operation for an encrypted volume is only supported when moving the volume to another aggregate. |
| key_id | string | The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key. |
| key_manager_attribute | string | Specifies an additional key manager attribute that is an identifier-value pair, separated by '='. For example, CRN=unique-value. This parameter is required when using the POST method and an IBM Key Lore key manager is configured on the SVM. |
| rekey | boolean | If set to 'true', re-encrypts the volume with a new key. Valid in PATCH. |
| state | string | Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one. |

| Name | Type | Description |
|--------|------------------------|---|
| status | status | |
| type | string | Volume encryption type. none ‐ The volume is a plain-text one. volume ‐ The volume is encrypted with NVE (NetApp Volume Encryption). aggregate ‐ The volume is encrypted with NAE (NetApp Aggregate Encryption). |

error_state

| Name | Type | Description |
|-----------------|---------|---|
| has_bad_blocks | boolean | Indicates whether the volume has any corrupt data blocks. If the damaged data block is accessed, an IO error, such as EIO for NFS or STATUS_FILE_CORRUPT for CIFS, is returned. |
| is_inconsistent | boolean | Indicates whether the file system has any inconsistencies. true ‐ File system is inconsistent. false ‐ File system 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. |

flash_pool

| Name | Type | Description |
|--------------------------|--------|---|
| cache_eligibility | string | If this parameter is specified, the command displays information only about the volume or volumes with the specified Flash Pool caching attributes. |
| cache_retention_priority | string | If this parameter is specified, the command displays the volumes that match the specified cache retention priority policy. A cache retention priority defines how long the blocks of a volume will be cached in the Flash Pool once they become cold. |

| Name | Type | Description |
|----------------|--------|---|
| caching_policy | string | This optionally specifies the caching policy to apply to the volume. A caching policy defines how the system caches a volume's data in Flash Cache modules. If a caching policy is not assigned to a volume, the system uses the caching policy that is assigned to the containing SVM. If a caching policy is not assigned to the containing SVM, the system uses the default cluster-wide policy. |

flexgroup

| Name | Type | Description |
|------|--------|---|
| name | string | Name of the FlexGroup volume that the constituent is part of. |
| uuid | string | Unique identifier for the FlexGroup volume that the constituent is part of. |

guarantee

| Name | Type | Description |
|---------|---------|---|
| honored | boolean | Is the space guarantee of this volume honored in the aggregate? |
| type | string | The type of space guarantee of this volume in the aggregate. |

iops

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

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

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

latency

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

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

cloud

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.

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

| Name | Type | Description |
|-----------|---------|---|
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

flexcache

Performance number for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
|--------------------|---------|--|
| bandwidth_savings | integer | Bandwidth savings denoting the amount of data served locally by the cache, in bytes. |
| cache_miss_percent | integer | Cache miss percentage. |

| Name | Type | Description |
|-----------|--------|---|
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

throughput

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

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

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

metric

Performance numbers, such as IOPS, latency and throughput.

| Name | Type | Description |
|-----------|---------------------------|--|
| _links | _links | |
| cloud | cloud | Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| flexcache | flexcache | Performance number for FlexCache used to measure cache effectiveness. |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
|------------|----------------------------|---|
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

movement

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

| Name | Type | Description |
|-----------------------|---------------------------------------|---|
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| start_time | string | Start time of volume move. |

| Name | Type | Description |
|----------------|--------|---|
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

export_policy

Export Policy

| Name | Type | Description |
|--------|------------------------|-------------|
| _links | _links | |
| id | integer | |
| name | string | |

junction_parent

| Name | Type | Description |
|--------|------------------------|--|
| _links | _links | |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |
| uuid | string | Unique identifier for the parent volume. |

nas

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

| Name | Type | Description |
|------------------|---------|---|
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth digit), first digit is assumed to be zero. |

policy

| Name | Type | Description |
|------------------------|------------------------|---|
| _links | _links | |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. |

| Name | Type | Description |
|---------------------|---------|---|
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

QoS information

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

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

| Name | Type | Description |
|---------|---------|--|
| enabled | boolean | This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVols or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVols or FlexGroup volumes when the quota state is "off". |
| state | string | Quota state of the volume |

retention

| Name | Type | Description |
|---------|--------|--|
| default | string | <p>Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period.</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 | Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none". |
| compliance_clock_time | string | This is the volume compliance clock time which is used to manage the SnapLock objects in the volume. |
| expiry_time | string | Expiry time of the volume. |
| is_audit_log | boolean | Indicates if this volume has been configured as SnapLock audit log volume for the SVM . |
| litigation_count | integer | Litigation count indicates the number of active legal-holds on the volume. |

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

destinations

| Name | Type | Description |
|----------|---------|---|
| is_cloud | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to a cloud destination. |
| is_ontap | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to an ONTAP destination. <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.9 |

snapmirror

Specifies attributes for SnapMirror protection.

| Name | Type | Description |
|--------------|------------------------------|---|
| destinations | destinations | |
| is_protected | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data. |

snapshot_policy

This is a reference to the Snapshot copy policy.

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

logical_space

| Name | Type | Description |
|-------------|---------|---|
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |
| used | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), in bytes. |
| used_by_afs | integer | The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes. |

| Name | Type | Description |
|-------------------|---------|--|
| used_by_snapshots | integer | Size that is logically used across all Snapshot copies in the volume, in bytes. |
| used_percent | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), as a percentage. |

snapshot

| Name | Type | Description |
|--------------------|---------|---|
| autodelete_enabled | boolean | Specifies whether Snapshot copy autodelete is currently enabled on this volume. |
| autodelete_trigger | string | Specifies when the system should trigger an autodelete of Snapshot copies. When set to <i>volume</i> , autodelete is triggered based on volume fullness. When set to <i>snap_reserve</i> , autodelete is triggered based on Snapshot reserve fullness. The default value is <i>volume</i> . |
| reserve_available | integer | Size available for Snapshot copies within the Snapshot copy reserve, in bytes. |
| reserve_percent | integer | The space that has been set aside as a reserve for Snapshot copy usage, in percent. |
| reserve_size | integer | Size in the volume that has been set aside as a reserve for Snapshot copy usage, in bytes. |
| space_used_percent | integer | Percentage of snapshot reserve size that has been used. |
| used | integer | The total space used by Snapshot copies in the volume, in bytes. |

space

| Name | Type | Description |
|---|---------|---|
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| 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. |
| block_storage_inactive_user_data_percent | integer | Percentage of size that is physically used in the performance tier of the volume. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |
| cross_volume_dedupe_metafiles_footprint | integer | Cross volume deduplication metadata footprint, in bytes. |
| cross_volume_dedupe_metafiles_temporary_footprint | integer | Cross volume temporary deduplication metadata footprint, in bytes. |
| dedupe_metafiles_footprint | integer | Deduplication metadata footprint, in bytes. |
| dedupe_metafiles_temporary_footprint | integer | Temporary deduplication metadata footprint, in bytes. |
| delayed_free_footprint | integer | Delayed free blocks footprint, in bytes. |
| expected_available | integer | Size that should be available for the volume, irrespective of available size in the aggregate, in bytes. |

| Name | Type | Description |
|-------------------------------|-------------------------------|---|
| file_operation_metadata | integer | File operation metadata footprint, in bytes. |
| filesystem_size | integer | Total usable size of the volume, in bytes. |
| filesystem_size_fixed | boolean | Specifies whether the file system is to remain of the same size when set to true or to grow when set to false. This option is automatically set to true when a volume becomes SnapMirrored. |
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| fractional_reserve | integer | Used to change the amount of space reserved for overwrites of reserved objects in a volume. |
| full_threshold_percent | integer | Volume full threshold percentage at which EMS warnings can be sent. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| nearly_full_threshold_percent | integer | Volume nearly full threshold percentage at which EMS warnings can be sent. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| overwrite_reserve | integer | Reserved space for overwrites, in bytes. |
| overwrite_reserve_used | integer | Overwrite logical reserve space used, in bytes. |

| Name | Type | Description |
|----------------------------------|--------------------------|---|
| percent_used | integer | Percentage of the volume size that is used. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| physical_used | integer | Size that is physically used in the volume, in bytes. |
| physical_used_percent | integer | Size that is physically used in the volume, as a percentage. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| size_available_for_snapshots | integer | Available space for Snapshot copies from snap-reserve, in bytes. |
| snapmirror_destination_footprint | integer | SnapMirror destination footprint, in bytes. |
| snapshot | snapshot | |
| snapshot_reserve_unusable | integer | Snapshot reserve that is not available for Snapshot copy creation, in bytes. |
| snapshot_spill | integer | Space used by the snapshot copies beyond the snap-reserve, in bytes. |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| user_data | integer | User data, in bytes. |

| Name | Type | Description |
|----------------------------|---------|---|
| volume_guarantee_footprint | integer | Space reserved for future writes in the volume, in bytes. |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

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

latency_raw

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

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

cloud

These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.

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

flexcache_raw

Performance numbers for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
|-------------------------|---------|---|
| cache_miss_blocks | integer | Blocks retrieved from origin in case of a cache miss. This can be divided by the raw client_requested_blocks and multiplied by 100 to calculate the cache miss percentage. |
| client_requested_blocks | integer | Total blocks requested by the client. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

throughput_raw

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

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

statistics

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

| Name | Type | Description |
|---------------|-------------------------------|--|
| cloud | cloud | These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools. |
| flexcache_raw | flexcache_raw | Performance numbers for FlexCache used to measure cache effectiveness. |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |

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

svm

SVM containing the volume. Required on POST.

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

tiering

| Name | Type | Description |
|------------------|---------------|--|
| min_cooling_days | integer | <p>This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. Note that this parameter is only used for tiering purposes and does not affect the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data. This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy.</p> |
| object_tags | array[string] | <p>This parameter specifies tags of a volume for objects stored on a FabricPool-enabled aggregate. Each tag is a key,value pair and should be in the format "key=value".</p> |

| Name | Type | Description |
|--------|--------|---|
| policy | string | <p>Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH. all &dash; This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks. auto &dash; This policy allows tiering of both snapshot and active file system user data to the cloud store none &dash; Volume blocks will not be tiered to the cloud store. snapshot_only &dash; This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days.</p> |

| Name | Type | Description |
|-----------|---------|---|
| supported | boolean | This parameter specifies whether or not FabricPools are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only FabricPool aggregates are used if this parameter is set to true and only non FabricPool aggregates are used if this parameter is set to false. Tiering support for a FlexGroup can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only does the volume need to support tiering by using FabricPools, the tiering "policy" must not be 'none'. A volume that uses FabricPools but has a tiering "policy" of 'none' supports tiering, but will not tier any data. |

error_arguments

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

error

| Name | Type | 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}

Introduced In: 9.6

Updates the attributes of a volume. For movement, use the "validate_only" field on the request to validate but not perform the operation. The PATCH API can be used to enable or disable quotas for a FlexVol or a FlexGroup volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline removes its junction path.

A PATCH request for volume encryption performs conversion/rekey operations asynchronously. You can retrieve the conversion/rekey progress details by calling a GET request on the corresponding volume endpoint.

Optional properties

- `queue_for_encryption` - Queue volumes for encryption when `encryption.enabled=true`. If this option is not provided or is false, conversion of volumes starts immediately. When there are volumes in the queue and less than four encryptions are running, volumes are encrypted in the order in which they are queued.
- `encryption.action` - You can pause an ongoing rekey/conversion operation or resume a paused rekey/conversion operation using this field. The following actions are supported for this field:
 - `conversion_pause` - Pause an encryption conversion operation currently in progress
 - `conversion_resume` - Resume a paused encryption conversion operation
 - `rekey_pause` - Pause an encryption rekey operation currently in progress
 - `rekey_resume` - Resume a paused encryption rekey operation

Related ONTAP commands

- `volume unmount`
- `volume mount`
- `volume online`
- `volume offline`
- `volume modify`
- `volume clone modify`
- `volume efficiency modify`
- `volume quota on`
- `volume quota off`
- `volume snaplock modify`
- `volume encryption conversion start`
- `volume encryption rekey start`
- `security anti-ransomware volume enable`
- `security anti-ransomware volume disable`
- `security anti-ransomware volume dry-run`

- security anti-ransomware volume pause
- security anti-ransomware volume resume

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. |
| restore_to.path | string | query | False | Path to the file which is restored from the Snapshot copy. • Introduced in: 9.10 |
| restore_to.restore_path | string | query | False | Specifies the destination location inside the volume where the file is restored. • Introduced in: 9.10 |
| restore_to.start_byte | number | query | False | Starting byte offset of the source file, in multiples of 4096. • Introduced in: 9.10 |

| Name | Type | In | Required | Description |
|-----------------------|---------|-------|----------|--|
| restore_to.byte_count | number | query | False | Number of bytes to restore from the source file, in multiples of 4096. <ul style="list-style-type: none"> • Introduced in: 9.10 |
| preserve_lun_ids | boolean | query | False | Specifies whether LUN IDs need to be preserved during a Snapshot copy restore operation. <ul style="list-style-type: none"> • Introduced in: 9.10 • Default value: |

| Name | Type | In | Required | Description |
|-----------------------------------|---------|-------|----------|---|
| nvfail | string | query | False | <p>When this option is "on", the filer performs additional work at boot time if it finds that there has been any potential data loss due to an NVRAM failure. In such situations, it causes the invalidation of all NFS file handles on all volumes affected by the problem so that client-side users are forced to remount the affected file system (and thus not continue to use potentially incorrect data). It is also possible to specify a set of files per volume that are renamed in such cases. The filer sends error messages to the console whenever such problems are found.</p> <ul style="list-style-type: none"> • Introduced in: 9.10 • enum: ["off", "on"] |
| snapshot_directory_access_enabled | boolean | query | False | <p>This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume.</p> <ul style="list-style-type: none"> • Introduced in: 9.10 • Default value: 1 |

| 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. This is limited to two new resources per available aggregate. • Default value: 1 • enum: ["use_existing_resources", "add_new_resources"] |

| Name | Type | In | Required | Description |
|----------------------------------|--------|-------|----------|--|
| scheduled_snapshot_naming_scheme | string | query | False | <p>Naming Scheme for automatic Snapshot copies:</p> <ul style="list-style-type: none"> • create_time - Automatic Snapshot copies are saved as per the start of their current date and time. • ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> • Introduced in: 9.10 • Default value: 1 • enum: ["create_time", "ordinal"] <p></scheduled_frequency></p> |

| Name | Type | In | Required | Description |
|---------------------------------|---------|-------|----------|--|
| clone.match_parent_storage_tier | boolean | query | False | <p>Specifies whether the FlexClone volume splits the data blocks by matching its parent storage tier. This option is applicable only if the tiering policy and the tiering minimum cooling days of the parent volume and the FlexClone volume are the same.</p> <ul style="list-style-type: none"> • Introduced in: 9.9 |
| return_timeout | integer | query | False | <p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0 |

| Name | Type | In | Required | Description |
|---------------|---------|-------|----------|---|
| validate_only | boolean | query | False | Validate the operation and its parameters, without actually performing the operation. |

Request Body

| Name | Type | Description |
|---------------------|-------------------------------------|---|
| _links | _links | |
| access_time_enabled | boolean | Indicates whether or not access time updates are enabled on the volume. |
| activity_tracking | activity_tracking | |
| aggregates | array[aggregates] | Aggregate hosting the volume. Required on POST. |
| analytics | analytics | |
| anti_ransomware | anti_ransomware | Anti-ransomware related information of the volume. |

| Name | Type | Description |
|-----------------------|-------------|---|
| anti_ransomware_state | string | <p>The Anti-ransomware state of the volume. If no "anti_ransomware_state" property is specified, the volume inherits the value from its parent SVM's "anti_ransomware_default_volume_state" property. If this value is "disabled", Anti-ransomware is disabled on the volume. If this value is "enabled", Anti-ransomware is enabled on the volume and alerts are raised if any suspect is detected for those volumes. If this value is "dry_run", Anti-ransomware is enabled in the dry-run or learning mode on the volume. The "dry_run" state is same as the "enabled" state except that the analytics data is used here for learning. No alerts are raised for any detections or violations. If this value is "paused", Anti-ransomware is paused on the volume. Additionally, three more states are available, which are only valid for GET. If this value is "disable_in_progress", Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. If this value is "enable_paused", Anti-ransomware is paused on the volume from its earlier enabled state. If this value is "dry_run_paused", Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. For POST, the valid Anti-ransomware states are only "disabled", "enabled" and "dry_run", whereas for PATCH, "paused" is also valid along with the three valid states for POST.</p> |
| application | application | |
| autosize | autosize | |
| clone | clone | |

| Name | Type | Description |
|------------------------|---------------------------------------|---|
| cloud_retrieval_policy | string | <p>This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved. "on_read" policy retrieves tiered data for all client driven data reads. "never" policy never retrieves tiered data. "promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved.</p> |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| consistency_group | consistency_group | Consistency group the volume is part of. |
| constituents | array[constituents] | |

| 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 volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list. |
| convert_unicode | boolean | Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients. |
| 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 | |
| flash_pool | flash_pool | |
| 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. |
| flexgroup | flexgroup | |
| guarantee | guarantee | |
| is_object_store | boolean | Specifies whether the volume is provisioned for an object store server. |

| Name | Type | Description |
|----------------------|--------------------------|--|
| is_svm_root | boolean | Specifies whether the volume is a root volume of the SVM it belongs to. |
| language | string | Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting. |
| max_dir_size | integer | Maximum directory size. This value sets maximum size, in bytes, to which a directory can grow. The default maximum directory size for FlexVol volumes is model-dependent, and optimized for the size of system memory. Before increasing the maximum directory size, involve technical support. |
| 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 |
| queue_for_encryption | boolean | Specifies whether the volume is queued for encryption. |

| Name | Type | Description |
|----------------------------------|---------------------------------|---|
| quota | quota | Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume. |
| scheduled_snapshot_naming_scheme | string | <p>Naming Scheme for automatic Snapshot copies:</p> <ul style="list-style-type: none"> • create_time - Automatic Snapshot copies are saved as per the start of their current date and time. • ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> |
| size | integer | Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size. |
| snaplock | snaplock | |
| snapmirror | snapmirror | Specifies attributes for SnapMirror protection. |
| snapshot_count | integer | Number of Snapshot copies in the volume. |
| snapshot_policy | snapshot_policy | This is a reference to the Snapshot copy policy. |
| space | space | |

| Name | Type | Description |
|------------|----------------------------|---|
| state | string | Volume state. Client access is supported only when volume is online and junctioned. Taking volume to offline or restricted state removes its junction path and blocks client access. When volume is in restricted state some operations like parity reconstruction and iron on commit are allowed. 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. |
| status | array[string] | Describes the current status of a volume. |

| Name | Type | Description |
|---------|---------|---|
| style | string | <p>The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol" you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. The style "flexgroup_constituent" is not supported when creating a volume. flexvol &dash; flexible volumes and FlexClone volumes flexgroup &dash; FlexGroup volumes flexgroup_constituent &dash; FlexGroup constituents.</p> |
| svm | svm | SVM containing the volume. Required on POST. |
| tiering | tiering | |
| type | string | <p>Type of the volume. rw &dash; read-write volume. dp &dash; data-protection volume. ls &dash; load-sharing <code>dp</code> volume. Valid in GET.</p> |

| 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 • Introduced in: 9.6 |

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "activity_tracking": {
    "state": "off",
    "unsupported_reason": {
      "code": "124518405",
      "message": "Volume activity tracking cannot be enabled on volumes
that contain LUNs."
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "analytics": {
    "scan_progress": 17,
    "state": "unknown",
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
    }
  },
  "anti_ransomware": {
    "attack_probability": "none",
    "attack_reports": {
      "_links": {
        "suspects": {
          "href": "/api/resourcelink"
        }
      }
    },
    "time": "2021-06-01T20:36:41+05:30"
  },
  "dry_run_start_time": "string",
  "space": {
```



```
    "snapshot_count": 0,
    "used": 0,
    "used_by_logs": 0,
    "used_by_snapshots": 0
  },
  "state": "disabled",
  "suspect_files": {
    "count": 0,
    "format": "string"
  }
},
"anti_ransomware_state": "disabled",
"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
  },
  "cloud_retrieval_policy": "default",
  "comment": "string",
  "consistency_group": {
    "name": "consistency_group_1",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "constituents": {
    "aggregates": {
      "name": "string",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "movement": {
      "cutover_window": 30,
      "destination_aggregate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "percent_complete": 0,
      "state": "replicating",
      "tiering_policy": "all"
    },
    "name": "string",
    "space": {
      "available": 0,
      "block_storage_inactive_user_data": 0,
      "capacity_tier_footprint": 0,
      "footprint": 0,
      "local_tier_footprint": 0,
      "logical_space": {
        "available": 0,
        "used_by_afs": 0
      },
      "metadata": 0,
      "over_provisioned": 0,
      "performance_tier_footprint": 0,
      "snapshot": {
        "used": 0
      }
    }
  }
}

```

```

    },
    "total_footprint": 0,
    "used": 0
  }
},
"create_time": "2018-06-04T19:00:00Z",
"efficiency": {
  "application_io_size": "8k",
  "compaction": "inline",
  "compression": "inline",
  "cross_volume_dedupe": "inline",
  "dedupe": "inline",
  "last_op_begin": "string",
  "last_op_end": "string",
  "last_op_err": "string",
  "last_op_size": 0,
  "last_op_state": "string",
  "op_state": "idle",
  "path": "string",
  "progress": "string",
  "schedule": "string",
  "state": "disabled",
  "storage_efficiency_mode": "default",
  "type": "regular"
},
"encryption": {
  "key_id": "string",
  "key_manager_attribute": "CRN=v1:bluemix:public:containers-
kubernetes:us-south:a/asdfghjkl1234:asdfghjkl1234:worker:kubernetes-
asdfghjkl-worker1",
  "state": "encrypted",
  "status": {
    "code": "string",
    "message": "string"
  },
  "type": "none"
},
"files": {
  "used": 0
},
"flash_pool": {
  "cache_eligibility": "read",
  "cache_retention_priority": "normal",
  "caching_policy": "none"
},
"flexcache_endpoint_type": "none",

```

```
"flexgroup": {
  "name": "my_flexgroup",
  "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
},
"guarantee": {
  "type": "volume"
},
"language": "ar",
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"cloud": {
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "timestamp": "2017-01-25T11:20:13Z"
},
"duration": "PT15S",
"flexcache": {
  "bandwidth_savings": 4096,
  "cache_miss_percent": 20,
  "duration": "PT1D",
  "status": "ok",
  "timestamp": "2017-01-25T11:20:13Z"
},
"iops": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"latency": {
  "read": 200,
  "total": 1000,
  "write": 100
}
```

```

    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": 0,
    "start_time": "2020-12-07T03:45:12-05:00",
    "state": "replicating",
    "tiering_policy": "all"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 100,
      "name": "default"
    },
    "junction_parent": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "vs1_root",
      "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
    },
    "path": "/user/my_volume",

```

```

    "security_style": "mixed",
    "unix_permissions": 755
  },
  "qos": {
    "policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    },
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"quota": {
  "state": "corrupt"
},
"scheduled_snapshot_naming_scheme": "create_time",
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04T19:00:00Z",
  "expiry_time": "Wed Sep 5 11:02:42 GMT 2018",
  "is_audit_log": 1,
  "litigation_count": 10,
  "privileged_delete": "enabled",
  "retention": {
    "default": "P30Y",
    "maximum": "P30Y",
    "minimum": "P30Y"
  },
  "type": "enterprise",
  "unspecified_retention_file_count": 10
},
"snapshot_count": 0,
"snapshot_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"name": "default",

```

```
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"space": {
  "available": 0,
  "block_storage_inactive_user_data": 0,
  "block_storage_inactive_user_data_percent": 0,
  "capacity_tier_footprint": 0,
  "cross_volume_dedupe_metafiles_footprint": 0,
  "cross_volume_dedupe_metafiles_temporary_footprint": 0,
  "dedupe_metafiles_footprint": 0,
  "dedupe_metafiles_temporary_footprint": 0,
  "delayed_free_footprint": 0,
  "file_operation_metadata": 0,
  "filesystem_size": 0,
  "footprint": 0,
  "local_tier_footprint": 0,
  "logical_space": {
    "available": 0,
    "used": 0,
    "used_by_afs": 0,
    "used_by_snapshots": 0,
    "used_percent": 0
  },
  "metadata": 0,
  "over_provisioned": 0,
  "overwrite_reserve": 0,
  "overwrite_reserve_used": 0,
  "percent_used": 0,
  "performance_tier_footprint": 0,
  "size_available_for_snapshots": 0,
  "snapmirror_destination_footprint": 0,
  "snapshot": {
    "autodelete_trigger": "volume",
    "reserve_available": 0,
    "reserve_size": 0,
    "space_used_percent": 0,
    "used": 0
  },
  "snapshot_reserve_unusable": 0,
  "snapshot_spill": 0,
  "total_footprint": 0,
  "used": 0,
  "user_data": 0,
  "volume_guarantee_footprint": 0
},
"state": "error",
```

```

"statistics": {
  "cloud": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "flexcache_raw": {
    "cache_miss_blocks": 10,
    "client_requested_blocks": 500,
    "status": "ok",
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "iops_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"status": {
},
"style": "flexvol",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
}

```



```

    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"tiering": {
  "object_tags": {
  },
  "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 |
|-------------------|---|
| 787141 | The specified "aggregates.name" and "aggregates.uuid" refer to different aggregates. |
| 917829 | Volume autosize grow threshold must be larger than autosize shrink threshold. |
| 917831 | Volume minimum autosize must be smaller than the maximum autosize. |
| 918193 | Cannot modify tiering min cooling days when vol move is in progress. |
| 918194 | Tiering min cooling days not supported for SVMDR. |
| 918195 | Tiering min cooling days not supported for non data volumes. |
| 918196 | Tiering min cooling days not allowed for the provided tiering policy. |
| 918248 | Specifying a value is not valid for initiating volume FlexClone split operation. |
| 918251 | Specifying a value is not valid for a Snapshot copy restore operation. |
| 918252 | specified "nas.path" is invalid. |
| 918265 | Volume is on the same aggregate. |
| 918266 | "movement.destination_aggregate" and "movement.state" are mutually exclusive, unless the state is "cutover-wait". |
| 918267 | The specified "movement.destination_aggregate" does not exist. |
| 918291 | Invalid volume cloud retrieval policy for the provided tiering policy. |
| 918292 | cloud retrieval policy not supported for non data volume. |
| 918293 | Cannot modify cloud retrieval policy when vol move is in progress. |
| 918521 | The volume maximum autosize must be smaller than or equal to the maximum volume size. |
| 918532 | The FlexClone match-parent-storage-tier option requires an effective cluster version of 9.9.1 or later. |
| 918533 | The FlexClone match-parent-storage-tier option not applicable for FlexClone volumes hosted on non-FabricPool storage. |
| 918534 | The tiering policy values are different for the FlexClone volume and its parent volume. The match-parent-storage-tier option cannot be set to true. |

| Error Code | Description |
|------------|--|
| 918535 | The tiering minimum cooling day values are different for the FlexClone volume and its parent volume. The match-parent-storage-tier option cannot be set to true. |
| 918537 | Could not get the FlexClone volume tiering policy or its parent volume tiering policy. Wait a minute and try again. |
| 918538 | The match-parent-storage-tier option is not supported for clone creation. |
| 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. |
| 13109198 | Resizing by adding new resources is only supported for FlexGroups. |
| 111411201 | File system analytics cannot be enabled on the target volume because of the specified reason. |
| 111411202 | File system analytics cannot be disabled on the target volume because of the specified reason. |
| 111411205 | File system analytics requires an effective cluster version of 9.8 or later. |
| 111411206 | The specified "analytics.state" is invalid. |
| 111411207 | File system analytics cannot be enabled on volumes that contain LUNs. |

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

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "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 | |

unsupported_reason

| Name | Type | Description |
|---------|--------|--|
| code | string | If volume activity tracking is not supported on the volume, this field provides an appropriate error code. |
| message | string | If volume activity tracking is not supported on the volume, this field provides an error message detailing why this is the case. |

activity_tracking

| Name | Type | Description |
|-------|--------|--|
| state | string | <p>Activity tracking state of the volume. If this value is "on", ONTAP collects top metrics information for the volume in real time. There is a slight impact to I/O performance in order to collect this information. If this value is "off", no activity tracking information is collected or available to view.</p> <ul style="list-style-type: none">• enum: ["off", "on"]• Introduced in: 9.10 |

| Name | Type | Description |
|--------------------|------------------------------------|--|
| supported | boolean | This field indicates whether or not volume activity tracking is supported on the volume. If volume activity tracking is not supported, the reason why is provided in the "activity_tracking.unsupported_reason" field. |
| unsupported_reason | unsupported_reason | |

aggregates

Aggregate

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

unsupported_reason

| Name | Type | Description |
|---------|--------|--|
| code | string | If file system analytics is not supported on the volume, this field provides the error code explaining why. |
| message | string | If file system analytics is not supported on the volume, this field provides the error message explaining why. |

analytics

| Name | Type | Description |
|---------------|---------|---|
| scan_progress | integer | Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is <code>initializing</code> . |

| Name | Type | Description |
|--------------------|------------------------------------|---|
| state | string | File system analytics state of the volume. If this value is "on", ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is "off", file system analytics information is not collected and not available to be viewed. If this value is "initializing", that means file system analytics was recently turned on, and the initialization scan to gather information all existing files and directories is currently running. If this value is 'unknown' that means there was an internal error when determining the file system analytics state for the volume. <ul style="list-style-type: none"> enum: ["unknown", "initializing", "off", "on"] Introduced in: 9.8 |
| supported | boolean | This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field. |
| unsupported_reason | unsupported_reason | |

_links

| Name | Type | Description |
|----------|----------------------|-------------|
| suspects | href | |

anti_ransomware_attack_report

| Name | Type | Description |
|--------|------------------------|---|
| _links | _links | |
| time | string | Timestamp at which ransomware attack is observed. |

space

| Name | Type | Description |
|-------------------|---------|--|
| snapshot_count | integer | Total number of Anti-ransomware backup Snapshot copies. |
| used | integer | Total space in bytes used by the Anti-ransomware feature. |
| used_by_logs | integer | Space in bytes used by the Anti-ransomware analytics logs. |
| used_by_snapshots | integer | Space in bytes used by the Anti-ransomware backup Snapshot copies. |

suspect_files

| Name | Type | Description |
|--------|---------|---|
| count | integer | Total number of <code>suspect_files.format</code> files observed by the Anti-ransomware analytics engine on the volume. |
| format | string | File formats observed by the Anti-ransomware analytics engine on the volume. |

anti_ransomware

Anti-ransomware related information of the volume.

| Name | Type | Description |
|--------------------|--|--|
| attack_probability | string | Probability of a ransomware attack. <code>none</code> No files are suspected of ransomware activity. <code>low</code> A number of files are suspected of ransomware activity. <code>moderate</code> A moderate number of files are suspected of ransomware activity. <code>high</code> A large number of files are suspected of ransomware activity. |
| attack_reports | array[anti_ransomware_attack_report] | |

| Name | Type | Description |
|--------------------|--|--|
| dry_run_start_time | string | Time when Anti-ransomware monitoring <code>state</code> is set to dry-run value for starting evaluation mode. |
| space | space | |
| state | string | <p>Anti-ransomware state.</p> <p><code>disabled</code> Anti-ransomware monitoring is disabled on the volume. This is the default state in a POST operation.</p> <p><code>disable_in_progress</code> Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. Valid in GET operation.</p> <p><code>dry_run</code> Anti-ransomware monitoring is enabled in the evaluation mode.</p> <p><code>enabled</code> Anti-ransomware monitoring is active on the volume.</p> <p><code>paused</code> Anti-ransomware monitoring is paused on the volume.</p> <p><code>enable_paused</code> Anti-ransomware monitoring is paused on the volume from its earlier enabled state. Valid in GET operation.</p> <p><code>dry_run_paused</code> Anti-ransomware monitoring is paused on the volume from its earlier <code>dry_run</code> state. Valid in GET operation. For POST, the valid Anti-ransomware states are only <code>disabled</code>, <code>enabled</code> and <code>dry_run</code>, whereas for PATCH, <code>paused</code> is also valid along with the three valid states for POST.</p> |
| suspect_files | array[suspect_files] | |

application

| Name | Type | Description |
|------|--------|--|
| name | string | Name of the application to which the volume belongs. Available only when the volume is part of an application. |

| Name | Type | Description |
|------|--------|--|
| uuid | string | UUID of the application to which the volume belongs. Available only when the volume is part of an application. |

autosize

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

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

snapshot_reference

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

parent_svm

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

parent_volume

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

| Name | Type | Description |
|------|--------|--|
| uuid | string | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6 |

clone

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

consistency_group

Consistency group the volume is part of.

| Name | Type | Description |
|------|--------|--|
| name | string | The name of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |
| uuid | string | The UUID of the consistency group to which the volume belongs. Available only when the volume is part of a consistency group. If this volume belongs to a child consistency group, then this will be the UUID of the parent consistency group. |

aggregates

| Name | Type | Description |
|------|--------|--|
| name | string | Name of the aggregate hosting the FlexGroup Constituent. |
| uuid | string | Unique identifier for the aggregate. |

destination_aggregate

Aggregate

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

movement

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

| Name | Type | Description |
|-----------------------|---------------------------------------|---|
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

logical_space

| Name | Type | Description |
|-------------|---------|--|
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |

| Name | Type | Description |
|-------------|---------|---|
| 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 |
|----------------------------------|---------|---|
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| block_storage_inactive_user_data | integer | The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |

| Name | Type | Description |
|----------------------------|-------------------------------|--|
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| snapshot | snapshot | |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| used_percent | integer | The virtual space used (includes volume reserves) before storage efficiency, as a percent. |

constituents

| Name | Type | Description |
|------------|----------------------------|-------------|
| aggregates | aggregates | |

| 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 | FlexGroup Constituents name |
| space | space | |

policy

| Name | Type | Description |
|------|--------|---|
| name | string | Specifies the name of the efficiency policy. Valid for PATCH. |

efficiency

| Name | Type | Description |
|---------------------|--------|--|
| application_io_size | string | Block size to use by compression. Valid for POST. |
| 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. |

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

| 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. |
| last_op_begin | string | Last sis operation begin timestamp. |
| last_op_end | string | Last sis operation end timestamp. |
| last_op_err | string | Last sis operation error text. |
| last_op_size | integer | Last sis operation size. |
| last_op_state | string | Last sis operation state. |
| op_state | string | Sis status of the volume. |
| path | string | Absolute volume path of the volume. |
| policy | policy | |
| progress | string | Sis progress of the volume. |
| schedule | string | Schedule associated with volume. |

| Name | Type | Description |
|-------------------------|--------|--|
| state | string | Storage efficiency state of the volume. Currently, this field supports POST/PATCH only for RW (Read-Write) volumes on FSx for ONTAP. disabled ‐ All storage efficiency features are disabled. mixed ‐ Read-only field for FlexGroup volumes, storage efficiency is enabled on certain constituents and disabled on others. On FSx for ONTAP ‐ enabled ‐ All supported storage efficiency features for the volume are enabled. custom ‐ Read-only field currently only supported for the FSx for ONTAP, user-defined storage efficiency features are enabled. For other platforms ‐ enabled ‐ At least one storage efficiency feature for the volume is enabled. <ul style="list-style-type: none"> enum: ["disabled", "enabled", "mixed", "custom"] Introduced in: 9.9 |
| storage_efficiency_mode | string | Storage efficiency mode used by volume. This parameter is supported only on AFF platform. |
| type | string | Sis Type of the volume. |

status

| Name | Type | Description |
|---------|--------|-----------------------------------|
| code | string | Encryption progress message code. |
| message | string | Encryption progress message. |

encryption

| Name | Type | Description |
|-----------------------|---------|--|
| enabled | boolean | Creates an encrypted or an unencrypted volume. For POST, when set to 'true', a new key is generated and used to encrypt the given volume. In that case, the underlying SVM must be configured with the key manager. When set to 'false', the volume created will be unencrypted. For PATCH, when set to 'true', it encrypts an unencrypted volume. Specifying the parameter as 'false' in a PATCH operation for an encrypted volume is only supported when moving the volume to another aggregate. |
| key_id | string | The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key. |
| key_manager_attribute | string | Specifies an additional key manager attribute that is an identifier-value pair, separated by '='. For example, CRN=unique-value. This parameter is required when using the POST method and an IBM Key Lore key manager is configured on the SVM. |
| rekey | boolean | If set to 'true', re-encrypts the volume with a new key. Valid in PATCH. |
| state | string | Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one. |

| Name | Type | Description |
|--------|------------------------|---|
| status | status | |
| type | string | Volume encryption type. none ‐ The volume is a plain-text one. volume ‐ The volume is encrypted with NVE (NetApp Volume Encryption). aggregate ‐ The volume is encrypted with NAE (NetApp Aggregate Encryption). |

error_state

| Name | Type | Description |
|-----------------|---------|---|
| has_bad_blocks | boolean | Indicates whether the volume has any corrupt data blocks. If the damaged data block is accessed, an IO error, such as EIO for NFS or STATUS_FILE_CORRUPT for CIFS, is returned. |
| is_inconsistent | boolean | Indicates whether the file system has any inconsistencies. true ‐ File system is inconsistent. false ‐ File system 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. |

flash_pool

| Name | Type | Description |
|--------------------------|--------|---|
| cache_eligibility | string | If this parameter is specified, the command displays information only about the volume or volumes with the specified Flash Pool caching attributes. |
| cache_retention_priority | string | If this parameter is specified, the command displays the volumes that match the specified cache retention priority policy. A cache retention priority defines how long the blocks of a volume will be cached in the Flash Pool once they become cold. |

| Name | Type | Description |
|----------------|--------|---|
| caching_policy | string | This optionally specifies the caching policy to apply to the volume. A caching policy defines how the system caches a volume's data in Flash Cache modules. If a caching policy is not assigned to a volume, the system uses the caching policy that is assigned to the containing SVM. If a caching policy is not assigned to the containing SVM, the system uses the default cluster-wide policy. |

flexgroup

| Name | Type | Description |
|------|--------|---|
| name | string | Name of the FlexGroup volume that the constituent is part of. |
| uuid | string | Unique identifier for the FlexGroup volume that the constituent is part of. |

guarantee

| Name | Type | Description |
|---------|---------|---|
| honored | boolean | Is the space guarantee of this volume honored in the aggregate? |
| type | string | The type of space guarantee of this volume in the aggregate. |

iops

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

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

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

latency

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

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

cloud

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.

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

| Name | Type | Description |
|-----------|---------|---|
| latency | latency | The round trip latency in microseconds observed at the storage object. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

flexcache

Performance number for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
|--------------------|---------|--|
| bandwidth_savings | integer | Bandwidth savings denoting the amount of data served locally by the cache, in bytes. |
| cache_miss_percent | integer | Cache miss percentage. |

| Name | Type | Description |
|-----------|--------|---|
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

throughput

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

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

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

metric

Performance numbers, such as IOPS, latency and throughput.

| Name | Type | Description |
|------------------------|---------------------------|--|
| _links | _links | |
| cloud | cloud | Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools. |
| duration | string | The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations: |
| flexcache | flexcache | Performance number for FlexCache used to measure cache effectiveness. |
| iops | iops | The rate of I/O operations observed at the storage object. |
| latency | latency | The round trip latency in microseconds observed at the storage object. |

| Name | Type | Description |
|------------|----------------------------|---|
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| throughput | throughput | The rate of throughput bytes per second observed at the storage object. |
| timestamp | string | The timestamp of the performance data. |

movement

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

| Name | Type | Description |
|-----------------------|---------------------------------------|---|
| cutover_window | integer | Time window in seconds for cutover. The allowed range is between 30 to 300 seconds. |
| destination_aggregate | destination_aggregate | Aggregate |
| percent_complete | integer | Completion percentage |
| start_time | string | Start time of volume move. |

| Name | Type | Description |
|----------------|--------|---|
| state | string | State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress. |
| tiering_policy | string | Tiering policy for FabricPool |

export_policy

Export Policy

| Name | Type | Description |
|--------|------------------------|-------------|
| _links | _links | |
| id | integer | |
| name | string | |

junction_parent

| Name | Type | Description |
|--------|------------------------|--|
| _links | _links | |
| name | string | The name of the parent volume that contains the junction inode of this volume. The junction parent volume must belong to the same SVM that owns this volume. |
| uuid | string | Unique identifier for the parent volume. |

nas

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

| Name | Type | Description |
|------------------|---------|---|
| unix_permissions | integer | UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth digit), first digit is assumed to be zero. |

policy

| Name | Type | Description |
|------------------------|------------------------|---|
| _links | _links | |
| max_throughput_iops | integer | Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. |
| max_throughput_mbps | integer | Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. |
| min_throughput_iops | integer | Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH. |

| Name | Type | Description |
|---------------------|---------|---|
| min_throughput_mbps | integer | Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH. |
| name | string | The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH. |
| uuid | string | The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH. |

qos

QoS information

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

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

| Name | Type | Description |
|---------|---------|--|
| enabled | boolean | This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVols or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVols or FlexGroup volumes when the quota state is "off". |
| state | string | Quota state of the volume |

retention

| Name | Type | Description |
|---------|--------|--|
| default | string | <p>Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period.</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 | Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none". |
| compliance_clock_time | string | This is the volume compliance clock time which is used to manage the SnapLock objects in the volume. |
| expiry_time | string | Expiry time of the volume. |
| is_audit_log | boolean | Indicates if this volume has been configured as SnapLock audit log volume for the SVM . |
| litigation_count | integer | Litigation count indicates the number of active legal-holds on the volume. |

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

destinations

| Name | Type | Description |
|----------|---------|---|
| is_cloud | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to a cloud destination. |
| is_ontap | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to an ONTAP destination. <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.9 |

snapmirror

Specifies attributes for SnapMirror protection.

| Name | Type | Description |
|--------------|------------------------------|---|
| destinations | destinations | |
| is_protected | boolean | Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data. |

snapshot_policy

This is a reference to the Snapshot copy policy.

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

logical_space

| Name | Type | Description |
|-------------|---------|---|
| available | integer | The amount of space available in this volume with storage efficiency space considered used, in bytes. |
| enforcement | boolean | Specifies whether space accounting for operations on the volume is done along with storage efficiency. |
| reporting | boolean | Specifies whether space reporting on the volume is done along with storage efficiency. |
| used | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), in bytes. |
| used_by_afs | integer | The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes. |

| Name | Type | Description |
|-------------------|---------|--|
| used_by_snapshots | integer | Size that is logically used across all Snapshot copies in the volume, in bytes. |
| used_percent | integer | SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), as a percentage. |

snapshot

| Name | Type | Description |
|--------------------|---------|---|
| autodelete_enabled | boolean | Specifies whether Snapshot copy autodelete is currently enabled on this volume. |
| autodelete_trigger | string | Specifies when the system should trigger an autodelete of Snapshot copies. When set to <i>volume</i> , autodelete is triggered based on volume fullness. When set to <i>snap_reserve</i> , autodelete is triggered based on Snapshot reserve fullness. The default value is <i>volume</i> . |
| reserve_available | integer | Size available for Snapshot copies within the Snapshot copy reserve, in bytes. |
| reserve_percent | integer | The space that has been set aside as a reserve for Snapshot copy usage, in percent. |
| reserve_size | integer | Size in the volume that has been set aside as a reserve for Snapshot copy usage, in bytes. |
| space_used_percent | integer | Percentage of snapshot reserve size that has been used. |
| used | integer | The total space used by Snapshot copies in the volume, in bytes. |

space

| Name | Type | Description |
|---|---------|---|
| afs_total | integer | Total size of AFS, excluding snap-reserve, in bytes. |
| available | integer | The available space, in bytes. |
| available_percent | integer | The space available, as a percent. |
| 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. |
| block_storage_inactive_user_data_percent | integer | Percentage of size that is physically used in the performance tier of the volume. |
| capacity_tier_footprint | integer | Space used by capacity tier for this volume in the FabricPool aggregate, in bytes. |
| cross_volume_dedupe_metafiles_footprint | integer | Cross volume deduplication metadata footprint, in bytes. |
| cross_volume_dedupe_metafiles_temporary_footprint | integer | Cross volume temporary deduplication metadata footprint, in bytes. |
| dedupe_metafiles_footprint | integer | Deduplication metadata footprint, in bytes. |
| dedupe_metafiles_temporary_footprint | integer | Temporary deduplication metadata footprint, in bytes. |
| delayed_free_footprint | integer | Delayed free blocks footprint, in bytes. |
| expected_available | integer | Size that should be available for the volume, irrespective of available size in the aggregate, in bytes. |

| Name | Type | Description |
|-------------------------------|-------------------------------|---|
| file_operation_metadata | integer | File operation metadata footprint, in bytes. |
| filesystem_size | integer | Total usable size of the volume, in bytes. |
| filesystem_size_fixed | boolean | Specifies whether the file system is to remain of the same size when set to true or to grow when set to false. This option is automatically set to true when a volume becomes SnapMirrored. |
| footprint | integer | Data used for this volume in the aggregate, in bytes. |
| fractional_reserve | integer | Used to change the amount of space reserved for overwrites of reserved objects in a volume. |
| full_threshold_percent | integer | Volume full threshold percentage at which EMS warnings can be sent. |
| local_tier_footprint | integer | Space used by the local tier for this volume in the aggregate, in bytes. |
| logical_space | logical_space | |
| metadata | integer | Space used by the volume metadata in the aggregate, in bytes. |
| nearly_full_threshold_percent | integer | Volume nearly full threshold percentage at which EMS warnings can be sent. |
| over_provisioned | integer | The amount of space not available for this volume in the aggregate, in bytes. |
| overwrite_reserve | integer | Reserved space for overwrites, in bytes. |
| overwrite_reserve_used | integer | Overwrite logical reserve space used, in bytes. |

| Name | Type | Description |
|----------------------------------|--------------------------|---|
| percent_used | integer | Percentage of the volume size that is used. |
| performance_tier_footprint | integer | Space used by the performance tier for this volume in the FabricPool aggregate, in bytes. |
| physical_used | integer | Size that is physically used in the volume, in bytes. |
| physical_used_percent | integer | Size that is physically used in the volume, as a percentage. |
| size | integer | Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes. |
| size_available_for_snapshots | integer | Available space for Snapshot copies from snap-reserve, in bytes. |
| snapmirror_destination_footprint | integer | SnapMirror destination footprint, in bytes. |
| snapshot | snapshot | |
| snapshot_reserve_unusable | integer | Snapshot reserve that is not available for Snapshot copy creation, in bytes. |
| snapshot_spill | integer | Space used by the snapshot copies beyond the snap-reserve, in bytes. |
| total_footprint | integer | Data and metadata used for this volume in the aggregate, in bytes. |
| used | integer | The virtual space used (includes volume reserves) before storage efficiency, in bytes. |
| used_by_afs | integer | The space used by Active Filesystem, in bytes. |
| user_data | integer | User data, in bytes. |

| Name | Type | Description |
|----------------------------|---------|---|
| volume_guarantee_footprint | integer | Space reserved for future writes in the volume, in bytes. |

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

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

latency_raw

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

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

cloud

These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.

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

flexcache_raw

Performance numbers for FlexCache used to measure cache effectiveness.

| Name | Type | Description |
|-------------------------|---------|---|
| cache_miss_blocks | integer | Blocks retrieved from origin in case of a cache miss. This can be divided by the raw client_requested_blocks and multiplied by 100 to calculate the cache miss percentage. |
| client_requested_blocks | integer | Total blocks requested by the client. |
| status | string | Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data. |
| timestamp | string | The timestamp of the performance data. |

throughput_raw

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

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

statistics

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

| Name | Type | Description |
|---------------|-------------------------------|--|
| cloud | cloud | These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools. |
| flexcache_raw | flexcache_raw | Performance numbers for FlexCache used to measure cache effectiveness. |
| iops_raw | iops_raw | The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time. |
| latency_raw | latency_raw | The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation. |

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

svm

SVM containing the volume. Required on POST.

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

tiering

| Name | Type | Description |
|------------------|---------------|--|
| min_cooling_days | integer | <p>This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. Note that this parameter is only used for tiering purposes and does not affect the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data. This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy.</p> |
| object_tags | array[string] | <p>This parameter specifies tags of a volume for objects stored on a FabricPool-enabled aggregate. Each tag is a key,value pair and should be in the format "key=value".</p> |

| Name | Type | Description |
|--------|--------|--|
| policy | string | <p>Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.</p> <p>all &dash; This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.</p> <p>auto &dash; This policy allows tiering of both snapshot and active file system user data to the cloud store</p> <p>none &dash; Volume blocks will not be tiered to the cloud store.</p> <p>snapshot_only &dash; This policy allows tiering of only the volume Snapshot copies not associated with the active file system.</p> <p>The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days.</p> |

| Name | Type | Description |
|-----------|---------|---|
| supported | boolean | This parameter specifies whether or not FabricPools are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only FabricPool aggregates are used if this parameter is set to true and only non FabricPool aggregates are used if this parameter is set to false. Tiering support for a FlexGroup can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only does the volume need to support tiering by using FabricPools, the tiering "policy" must not be 'none'. A volume that uses FabricPools but has a tiering "policy" of 'none' supports tiering, but will not tier any data. |

volume

| Name | Type | Description |
|---------------------|-------------------------------------|---|
| _links | _links | |
| access_time_enabled | boolean | Indicates whether or not access time updates are enabled on the volume. |
| activity_tracking | activity_tracking | |
| aggregates | array[aggregates] | Aggregate hosting the volume. Required on POST. |
| analytics | analytics | |
| anti_ransomware | anti_ransomware | Anti-ransomware related information of the volume. |

| Name | Type | Description |
|-----------------------|-----------------------------|---|
| anti_ransomware_state | string | <p>The Anti-ransomware state of the volume. If no "anti_ransomware_state" property is specified, the volume inherits the value from its parent SVM's "anti_ransomware_default_volume_state" property. If this value is "disabled", Anti-ransomware is disabled on the volume. If this value is "enabled", Anti-ransomware is enabled on the volume and alerts are raised if any suspect is detected for those volumes. If this value is "dry_run", Anti-ransomware is enabled in the dry-run or learning mode on the volume. The "dry_run" state is same as the "enabled" state except that the analytics data is used here for learning. No alerts are raised for any detections or violations. If this value is "paused", Anti-ransomware is paused on the volume. Additionally, three more states are available, which are only valid for GET. If this value is "disable_in_progress", Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. If this value is "enable_paused", Anti-ransomware is paused on the volume from its earlier enabled state. If this value is "dry_run_paused", Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. For POST, the valid Anti-ransomware states are only "disabled", "enabled" and "dry_run", whereas for PATCH, "paused" is also valid along with the three valid states for POST.</p> |
| application | application | |
| autosize | autosize | |
| clone | clone | |

| Name | Type | Description |
|------------------------|---------------------------------------|---|
| cloud_retrieval_policy | string | <p>This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved. "on_read" policy retrieves tiered data for all client driven data reads. "never" policy never retrieves tiered data. "promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved.</p> |
| comment | string | A comment for the volume. Valid in POST or PATCH. |
| consistency_group | consistency_group | Consistency group the volume is part of. |
| constituents | array[constituents] | |

| 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 volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list. |
| convert_unicode | boolean | Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients. |
| 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 | |
| flash_pool | flash_pool | |
| 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. |
| flexgroup | flexgroup | |
| guarantee | guarantee | |

| Name | Type | Description |
|-----------------|--------------------------|--|
| is_object_store | boolean | Specifies whether the volume is provisioned for an object store server. |
| is_svm_root | boolean | Specifies whether the volume is a root volume of the SVM it belongs to. |
| language | string | Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting. |
| max_dir_size | integer | Maximum directory size. This value sets maximum size, in bytes, to which a directory can grow. The default maximum directory size for FlexVol volumes is model-dependent, and optimized for the size of system memory. Before increasing the maximum directory size, involve technical support. |
| 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 | |

| Name | Type | Description |
|----------------------------------|---------------------------------|---|
| qos | qos | QoS information |
| queue_for_encryption | boolean | Specifies whether the volume is queued for encryption. |
| quota | quota | Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume. |
| scheduled_snapshot_naming_scheme | string | <p>Naming Scheme for automatic Snapshot copies:</p> <ul style="list-style-type: none"> • create_time - Automatic Snapshot copies are saved as per the start of their current date and time. • ordinal - Latest automatic snapshot copy is saved as <scheduled_frequency>.0 and subsequent copies will follow the create_time naming convention.</scheduled_frequency> |
| size | integer | Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size. |
| snaplock | snaplock | |
| snapmirror | snapmirror | Specifies attributes for SnapMirror protection. |
| snapshot_count | integer | Number of Snapshot copies in the volume. |
| snapshot_policy | snapshot_policy | This is a reference to the Snapshot copy policy. |
| space | space | |

| Name | Type | Description |
|------------|----------------------------|---|
| state | string | Volume state. Client access is supported only when volume is online and junctioned. Taking volume to offline or restricted state removes its junction path and blocks client access. When volume is in restricted state some operations like parity reconstruction and iron on commit are allowed. 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. |
| status | array[string] | Describes the current status of a volume. |

| Name | Type | Description |
|---------|---------|---|
| style | string | <p>The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol" you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. The style "flexgroup_constituent" is not supported when creating a volume. flexvol &dash; flexible volumes and FlexClone volumes flexgroup &dash; FlexGroup volumes flexgroup_constituent &dash; FlexGroup constituents.</p> |
| svm | svm | SVM containing the volume. Required on POST. |
| tiering | tiering | |
| type | string | <p>Type of the volume. rw &dash; read-write volume. dp &dash; data-protection volume. ls &dash; load-sharing <code>dp</code> volume. Valid in GET.</p> |

| Name | Type | Description |
|-------------------------|---------|--|
| use_mirrored_aggregates | boolean | Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration. |
| uuid | string | <p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6 |

job_link

| Name | Type | Description |
|------------------------|------------------------|---|
| _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.