Manage storage volumes

REST API reference

NetApp
September 09, 2025

This PDF was generated from https://docs.netapp.com/us-en/ontap-restapi-9161/storage_volumes_endpoint_overview.html on September 09, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Manage storage volumes . . . 1
- Storage volumes endpoint overview . . . 1
  - Overview . . . 1
  - Storage efficiency . . . 1
  - Quotas . . . 1
  - File System Analytics . . . 1
  - QoS . . . 2
  - Performance monitoring . . . 2
  - Rebalancing . . . 2
  - Volume APIs . . . 2
  - Examples . . . 3
  - Volumes reported in the GET REST API . . . 7
  - Examples . . . 8
  - Updating the attributes of a volume . . . 17
  - Examples . . . 17
  - Add tiering object tags for a FlexVol volume . . . 30
  - Deleting a volume . . . 31
  - Example . . . 31
- Retrieve volumes . . . 33
  - Expensive properties . . . 33
  - Related ONTAP commands . . . 37
  - Parameters . . . 37
  - Response . . . 114
  - Error . . . 143
  - Definitions . . . 143
- Create a volume on an SVM and storage aggregates . . . 239
  - Platform Specifics . . . 239
  - Required properties . . . 239
  - Default property values . . . 240
  - Related ONTAP commands . . . 240
  - Parameters . . . 240
  - Request Body . . . 241
  - Response . . . 262
  - Response . . . 263
  - Error . . . 263
  - Definitions . . . 268
- Delete a volume . . . 355
  - Platform Specifics . . . 355
  - Optional parameters: . . . 355
  - Related ONTAP commands . . . 355
  - Parameters . . . 355
  - Response . . . 356
  - Response . . . 357

Error . . . 357
Definitions . . . 359
Retrieve a volume . . . 360
Expensive properties . . . 361
Related ONTAP commands . . . 363
Parameters . . . 364
Response . . . 364
Error . . . 403
Definitions . . . 403
Update volume attributes . . . 490
Platform Specifics . . . 490
Optional properties . . . 490
Related ONTAP commands . . . 490
Parameters . . . 491
Request Body . . . 498
Response . . . 518
Response . . . 519
Error . . . 519
Definitions . . . 527

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. It is possible to create a SnapLock volume by specifying SnapLock parameters.

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, compaction, policy-name, enabled, application_io_size, compression_type and storage_efficiency_mode 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. Application-io-size and compression-type decides type of compression behavior in the system. Storage efficiency mode decides if the system is to run in default/efficient mode. Detailed information about each field is available under efficiency object for storage efficiency fields.

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 (FSA) provide a quick method for obtaining information summarizing properties of all files within any directory tree of a volume. For more information on FSA, 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

If the `analytics.state` field is "initializing" or "initialization_paused" and is set to "off", the FSA initialization scan is cancelled. If FSA is turned on again, the initialization scan restarts.

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.

Rebalancing

Non-disruptive capacity rebalancing of a FlexGroup volume is configured by the `rebalancing.*` fields. If not explicitly set, default values are provided. To initiate a capacity rebalancing operation, `rebalancing.state` is set to 'starting'. The `rebalancing.max_runtime` can be optionally set, which is the maximum length of time you want the capacity rebalancing to run for. You can stop capacity rebalancing by setting `rebalancing.state` to 'stopping'. You can also modify the configurations `rebalancing.max_runtime`, `rebalancing.max_threshold`, `rebalancing.min_threshold`, `rebalancing.max_file_moves`, `rebalancing.min_file_size`, and `rebalancing.exclude_snapshots`. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current configuration values. Once the operation is started, changes to the configuration are not allowed, until the capacity rebalanding operations stops, either by exceeding their maximum runtime or by being stopped. To see runtime information about each constituent, for a running rebalancing operation, use the 'rebalancing.engine.*' fields.

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" : "vol1", "state" : "online", "type" :
"RW", "aggregates" : [{"name" : "aggr1"}, {"name" : "aggr2"},
{"name":"aggr3"}], "constituents_per_aggregate" : "1", "svm" : {"name" :
"vs1"}, "size" : "240MB", "encryption" : {"enabled" : "False"},
"efficiency" : {"compression" : "both"}, "autosize" : {"maximum" :
"500MB", "minimum" : "240MB"}}'

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

Creating a FlexGroup volume and specifying its properties using POST when the Performance_NAS license is installed.

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

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name" : "vol1", "state" : "online", "type" :
"RW", "svm" : {"name" : "vs1"}, "size" : "240TB", "encryption" :
{"enabled" : "False"}, "efficiency" : {"compression" : "both"}, "autosize"
: {"maximum" : "500TB", "minimum" : "240TB"}}'

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

Creating a FlexClone and specifying its properties using POST

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

# The call:
curl -X POST  "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name":"vol1_clone", "clone": {"parent_volume":
{"name": "vol1"}, "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=vol1_clone
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "c9ee0040-39a5-11e9-9b24-00a098439a83",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/c9ee0040-39a5-11e9-9b24-00a098439a83"
    }
  }
}
}
```

Volumes reported in the GET REST API

The following types of volumes are reported:

– RW, DP and LS 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": "vol1",
    "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"
}
}
```

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?flexgroup.uuid=fd87d06f-8876-11ec-94a3-
005056a7484f&is_constituent=true" -H "accept: application/hal+json"

# The response:
{
"records": [
  {
    "uuid": "fd877f7c-8876-11ec-94a3-005056a7484f",
    "name": "fg__0001",
    "flexgroup": {
      "uuid": "fd87d06f-8876-11ec-94a3-005056a7484f"
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/fd877f7c-8876-11ec-94a3-
005056a7484f?is_constituent=true"
```

```
      }
    }
  },
  {
    "uuid": "fea631d6-8876-11ec-94a3-005056a7484f",
    "name": "fg__0002",
    "flexgroup": {
      "uuid": "fd87d06f-8876-11ec-94a3-005056a7484f"
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/fea631d6-8876-11ec-94a3-
005056a7484f?is_constituent=true"
      }
    }
  },
  {
    "uuid": "ff38a34e-8876-11ec-94a3-005056a7484f",
    "name": "fg__0003",
    "flexgroup": {
      "uuid": "fd87d06f-8876-11ec-94a3-005056a7484f"
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/ff38a34e-8876-11ec-94a3-
005056a7484f?is_constituent=true"
      }
    }
  },
  {
    "uuid": "ffdbbd1f-8876-11ec-94a3-005056a7484f",
    "name": "fg__0004",
    "flexgroup": {
      "uuid": "fd87d06f-8876-11ec-94a3-005056a7484f"
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/ffdbbd1f-8876-11ec-94a3-
005056a7484f?is_constituent=true"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
```

```
    "href": "/api/storage/volumes?flexgroup.uuid=fd87d06f-8876-11ec-94a3-
005056a7484f&is_constituent=true"
  }
}
}
```

Retrieving the efficiency attributes of volume

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

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/5f098ebc-32c8-11eb-
8dde-005056ace228/?fields=efficiency" -H "accept: application/hal+json"

# The response:
{
"uuid": "5f098ebc-32c8-11eb-8dde-005056ace228",
"name": "vol1",
"efficiency": {
  "compression": "both",
  "dedupe": "background",
  "cross_volume_dedupe": "none",
  "compaction": "none",
  "schedule": "sun-sat@0",
  "svm": "vs0",
  "state": "enabled",
  "status": "idle",
  "type": "regular",
  "progress": "Idle for 00:10:37",
  "last_op_begin_timestamp": "Mon Nov 30 00:00:02 2020",
  "last_op_end_timestamp": "Mon Nov 30 00:00:03 2020",
  "last_op_state": "Success",
  "last_op_size": 0,
  "addr": "/vol/vol1",
  "policy": {
    "name": "-"
  }
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/5f098ebc-32c8-11eb-8dde-005056ace228"
   }
}
}
```

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"
    }
  }
}
}
```

Stopping a volume clone split operation on 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":"false"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 03 Nov 2021 15:10:04 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": "1f8b3673-3cb8-11ec-b89e-005056bb2cb5",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/1f8b3673-3cb8-11ec-b89e-005056bb2cb5"
    }
  }
}
}
```

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"
    }
  }
}
}
```

Starting non-disruptive volume capacity rebalancing for 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 '{"rebalancing":{"state":"starting",
"max_runtime":"PT6H"}}' -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"
    }
  }
}
}
```

Starting a scheduled non-disruptive volume capacity rebalancing for 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 '{"rebalancing":{"state":"starting",
"start_time":"2022-12-21T15:30:00-05:00"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 21 Dec 2022 15:25:00 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": "53b2a954-80d4-11ed-81d2-005056ac8ed0",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/53b2a954-80d4-11ed-81d2-005056ac8ed0"
    }
  }
}
}
```

Stopping non-disruptive volume capacity rebalancing OR scheduled rebalancing for a FlexGroup volume using PATCH. This works for scheduled or on-going rebalancing.

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

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"rebalancing":{"state":"stopping"}}' -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"
    }
  }
}
}
```

Modifying non-disruptive volume capacity rebalancing configurations for a FlexGroup volume

The following example shows how to use a PATCH request to modify non-disruptive volume capacity rebalancing configurations for a FlexGroup volume:

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

# The call:
curl -X PATCH  "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"rebalancing":{"start_time":"2023-03-18T15:30:00-
05:00","max_threshold":20, "min_threshold":5, "max_file_moves":15,
"min_file_size":"100MB", "exclude_snapshots":"false",
"max_runtime":"PT6H"}}' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Thu, 16 Mar 2023 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"
    }
  }
}
}
```

Retrieving non-disruptive volume capacity rebalancing engine runtime information for a FlexGroup volume

The following example shows how to use a GET request to retrieve non-disruptive volume capacity rebalancing engine runtime information for a FlexGroup volume:

```
# The API:
/api/storage/volumes?is_constituent=true

# The call:
curl -X GET  "https://<mgmt-
ip>/api/storage/volumes?fields=rebalancing.engine&is_constituent=true&flex
group.uuid=d0c3359c-5448-4a9b-a077-e3295a7e9057" -H "accept:
application/hal+json"

# The response:
```

```
{
"records": [
  {
    "uuid": "2b32fdf1-b916-11ec-b103-005056a79638",
    "name": "fg__0001",
    "flexgroup": {
      "uuid": "2b3323db-b916-11ec-b103-005056a79638"
    },
    "rebalancing": {
      "engine": {
        "scanner": {
          "files_scanned": 3522915,
          "files_skipped": {
            "too_small": 3812,
            "too_large": 199,
            "fast_truncate": 22,
            "in_snapshot": 77499,
            "efficiency_blocks": 1823,
            "efficiency_percent": 355,
            "incompatible": 9377,
            "metadata": 85449,
            "remote_cache": 1912,
            "write_fenced": 28,
            "on_demand_destination": 87,
            "footprint_invalid": 12,
            "other": 336
          },
          "blocks_scanned": 1542675000,
          "blocks_skipped": {
            "too_small": 8744000,
            "too_large": 865000,
            "fast_truncate": 54000,
            "in_snapshot": 7749000,
            "efficiency_blocks": 1472000,
            "efficiency_percent": 366000,
            "incompatible": 2287000,
            "metadata": 85673000,
            "remote_cache": 9914000,
            "write_fenced": 19000,
            "on_demand_destination": 66000,
            "footprint_invalid": 98000,
            "other":187000
          }
        },
        "movement": {
          "file_moves_started": 9833,
```

```
          "most_recent_start_time": "2022-02-15T12:56:07-05:00",
          "last_error": {
            "time": "2022-02-15T09:09:27-05:00",
            "file_id": 88,
            "destination": 1089,
            "code": 60
          }
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/2b32fdf1-b916-11ec-b103-
005056a79638?is_constituent=true"
      }
    }
  },
  {
    "uuid": "2cc5da55-b916-11ec-b103-005056a79638",
    "name": "fg__0002",
    "flexgroup": {
      "uuid": "2b3323db-b916-11ec-b103-005056a79638"
    },
    "rebalancing": {
      "engine": {
        "scanner": {
          "files_scanned": 3522915,
          "files_skipped": {
            "too_small": 3812,
            "too_large": 188,
            "fast_truncate": 25,
            "in_snapshot": 77499,
            "efficiency_blocks": 1823,
            "efficiency_percent": 355,
            "incompatible": 9377,
            "metadata": 85449,
            "remote_cache": 1912,
            "write_fenced": 28,
            "on_demand_destination": 87,
            "footprint_invalid": 12,
            "other": 336
          },
          "blocks_scanned": 1542675000,
          "blocks_skipped": {
            "too_small": 8744000,
            "too_large": 865000,
```

```
            "fast_truncate": 54000,
            "in_snapshot": 7749000,
            "efficiency_blocks": 1472000,
            "efficiency_percent": 366000,
            "incompatible": 2287000,
            "metadata": 85673000,
            "remote_cache": 9914000,
            "write_fenced": 19000,
            "on_demand_destination": 66000,
            "footprint_invalid": 98000,
            "other":187000
          }
        },
        "movement": {
          "file_moves_started": 9833,
          "most_recent_start_time": "2022-02-15T12:56:07-05:00",
          "last_error": {
            "time": "2022-02-15T08:09:27-05:00",
            "file_id": 88,
            "destination": 1089,
            "code": 60
          }
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/2cc5da55-b916-11ec-b103-
005056a79638?is_constituent=true"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href":
"/api/storage/volumes?fields=rebalancing.engine&is_constituent=true&flexgr
oup.uuid=d0c3359c-5448-4a9b-a077-e3295a7e9057"
  }
}
}
```

Add tiering object tags for a FlexVol volume

The following example shows how to use a PATCH request to add tiering object tags for a FlexVol volume:

```
# 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/{uuid}

# 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"
    }
  }
}
}
```

Deleting a volume and bypassing the recovery queue

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

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

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 09 Feb 2022 09:59:55 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Content-Length: 189
Content-Type: application/hal+json
{
"job": {
  "uuid": "08757020-898f-11ec-b367-005056bb7353",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/08757020-898f-11ec-b367-005056bb7353"
    }
  }
}
}
```

Retrieve volumes

GET `/storage/volumes`

Introduced In: 9.6

Retrieves volumes.

Expensive properties

There is an added computational 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`
- `aggressive_readahead_mode`

- 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.space.large_size_enabled
- constituents.space.max_size
- constituents.space.total_metadata
- constituents.space.total_metadata_footprint
- 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
- asynchronous_directory_delete.*
- rebalancing.*
- metric.*
- cloud_write_enabled

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
- volume rebalance show
- security anti-ransomware volume show
- security anti-ransomware volume space show
- volume file async-delete client 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. • Default value: • Introduced in: 9.10
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.name	string	query	False	Filter by constituents.name • Introduced in: 9.9
constituents.space.available_percent	integer	query	False	Filter by constituents.space.available_percent • Introduced in: 9.9
constituents.space.over_provisioned	integer	query	False	Filter by constituents.space.over_provisioned • Introduced in: 9.9
constituents.space.local_tier_footprint	integer	query	False	Filter by constituents.space.local_tier_footprint • Introduced in: 9.9

Name	Type	In	Required	Description
constituents.space.afs_total	integer	query	False	Filter by constituents.space.afs_total • 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
constituents.space.snapshot.used	integer	query	False	Filter by constituents.space.snapshot.used • 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.total_footprint	integer	query	False	Filter by constituents.space.total_footprint • Introduced in: 9.9

Name	Type	In	Required	Description
constituents.space.large_size_enabled	boolean	query	False	Filter by constituents.space.large_size_enabled • Introduced in: 9.12
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.enforcement	boolean	query	False	Filter by constituents.space.logical_space.enforcement • 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.block_storage_inactive_user_data	integer	query	False	Filter by constituents.space.block_storage_inactive_user_data • Introduced in: 9.9

Name	Type	In	Required	Description
constituents.space.used	integer	query	False	Filter by constituents.space.used • Introduced in: 9.9
constituents.space.total_metadata_footprint	integer	query	False	Filter by constituents.space.total_metadata_footprint • Introduced in: 9.15
constituents.space.capacity_tier_footprint	integer	query	False	Filter by constituents.space.capacity_tier_footprint • Introduced in: 9.9
constituents.space.total_metadata	integer	query	False	Filter by constituents.space.total_metadata • Introduced in: 9.15
constituents.space.footprint	integer	query	False	Filter by constituents.space.footprint • Introduced in: 9.9
constituents.space.used_percent	integer	query	False	Filter by constituents.space.used_percent • Introduced in: 9.10

Name	Type	In	Required	Description
constituents.space.performance_tier_footprint	integer	query	False	Filter by constituents.space.performance_tier_footprint • Introduced in: 9.9
constituents.space.used_by_afs	integer	query	False	Filter by constituents.space.used_by_afs • Introduced in: 9.9
constituents.space.max_size	string	query	False	Filter by constituents.space.max_size • Introduced in: 9.14
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
constituents.movement.percent_complete	integer	query	False	Filter by constituents.movement.percent_complete • Introduced in: 9.9

Name	Type	In	Required	Description
constituents.movement.state	string	query	False	Filter by constituents.movement.state • Introduced in: 9.9
constituents.movement.cutover_window	integer	query	False	Filter by constituents.movement.cutover_window • Introduced in: 9.9
constituents.aggregates.name	string	query	False	Filter by constituents.aggregates.name • Introduced in: 9.9
constituents.aggregates.uuid	string	query	False	Filter by constituents.aggregates.uuid • Introduced in: 9.9
convert_unicode	boolean	query	False	Filter by convert_unicode • Introduced in: 9.10
style	string	query	False	Filter by style
language	string	query	False	Filter by language
cloud_write_enabled	boolean	query	False	Filter by cloud_write_enabled • Introduced in: 9.13
rebalancing.state	string	query	False	Filter by rebalancing.state • Introduced in: 9.11

Name	Type	In	Required	Description
rebalancing.stop_time	string	query	False	Filter by rebalancing.stop_time • Introduced in: 9.11
rebalancing.exclude_snapshots	boolean	query	False	Filter by rebalancing.exclude_snapshots • Introduced in: 9.11
rebalancing.max_runtime	string	query	False	Filter by rebalancing.max_runtime • Introduced in: 9.11
rebalancing.engine.movement.most_recent_start_time	string	query	False	Filter by rebalancing.engine.movement.most_recent_start_time • Introduced in: 9.12
rebalancing.engine.movement.last_error.code	integer	query	False	Filter by rebalancing.engine.movement.last_error.code • Introduced in: 9.12
rebalancing.engine.movement.last_error.file_id	integer	query	False	Filter by rebalancing.engine.movement.last_error.file_id • Introduced in: 9.12

Name	Type	In	Required	Description
rebalancing.engine.movement.last_error.time	string	query	False	Filter by rebalancing.engine.movement.last_error.time • Introduced in: 9.12
rebalancing.engine.movement.last_error.destination	integer	query	False	Filter by rebalancing.engine.movement.last_error.destination • Introduced in: 9.12
rebalancing.engine.movement.file_moves_started	integer	query	False	Filter by rebalancing.engine.movement.file_moves_started • Introduced in: 9.12
rebalancing.engine.scanner.blocks_skipped.in_snapshot	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.in_snapshot • Introduced in: 9.12
rebalancing.engine.scanner.blocks_skipped.too_large	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.too_large • Introduced in: 9.12
rebalancing.engine.scanner.blocks_skipped.fast_truncate	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.fast_truncate • Introduced in: 9.12

Name	Type	In	Required	Description
rebalancing.engine.scanner.blocks_skipped.footprint_invalid	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.footprint_invalid • Introduced in: 9.12
rebalancing.engine.scanner.blocks_skipped.incompatible	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.incompatible • Introduced in: 9.12
rebalancing.engine.scanner.blocks_skipped.efficiency_blocks	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.efficiency_blocks • Introduced in: 9.12
rebalancing.engine.scanner.blocks_skipped.other	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.other • Introduced in: 9.12
rebalancing.engine.scanner.blocks_skipped.efficiency_percent	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.efficiency_percent • Introduced in: 9.12
rebalancing.engine.scanner.blocks_skipped.metadata	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.metadata • Introduced in: 9.12

Name	Type	In	Required	Description
rebalancing.engine.scanner.blocks_skipped.on_demand_destination	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.on_demand_destination • Introduced in: 9.12
rebalancing.engine.scanner.blocks_skipped.write_fenced	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.write_fenced • Introduced in: 9.12
rebalancing.engine.scanner.blocks_skipped.remote_cache	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.remote_cache • Introduced in: 9.12
rebalancing.engine.scanner.blocks_skipped.too_small	integer	query	False	Filter by rebalancing.engine.scanner.blocks_skipped.too_small • Introduced in: 9.12
rebalancing.engine.scanner.files_skipped.in_snapshot	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.in_snapshot • Introduced in: 9.12
rebalancing.engine.scanner.files_skipped.fast_truncate	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.fast_truncate • Introduced in: 9.12

Name	Type	In	Required	Description
rebalancing.engine.scanner.files_skipped.footprint_invalid	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.footprint_invalid • Introduced in: 9.12
rebalancing.engine.scanner.files_skipped.too_large	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.too_large • Introduced in: 9.12
rebalancing.engine.scanner.files_skipped.other	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.other • Introduced in: 9.12
rebalancing.engine.scanner.files_skipped.efficiency_percent	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.efficiency_percent • Introduced in: 9.12
rebalancing.engine.scanner.files_skipped.incompatible	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.incompatible • Introduced in: 9.12
rebalancing.engine.scanner.files_skipped.efficiency_blocks	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.efficiency_blocks • Introduced in: 9.12

Name	Type	In	Required	Description
rebalancing.engine.scanner.files_skipped.write_fenced	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.write_fenced • Introduced in: 9.12
rebalancing.engine.scanner.files_skipped.too_small	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.too_small • Introduced in: 9.12
rebalancing.engine.scanner.files_skipped.remote_cache	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.remote_cache • Introduced in: 9.12
rebalancing.engine.scanner.files_skipped.metadata	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.metadata • Introduced in: 9.12
rebalancing.engine.scanner.files_skipped.on_demand_destination	integer	query	False	Filter by rebalancing.engine.scanner.files_skipped.on_demand_destination • Introduced in: 9.12
rebalancing.engine.scanner.files_scanned	integer	query	False	Filter by rebalancing.engine.scanner.files_scanned • Introduced in: 9.12

Name	Type	In	Required	Description
rebalancing.engine.scanner.blocks_scanned	integer	query	False	Filter by rebalancing.engine.scanner.blocks_scanned • Introduced in: 9.12
rebalancing.start_time	string	query	False	Filter by rebalancing.start_time • Introduced in: 9.11
rebalancing.imbalance_percent	integer	query	False	Filter by rebalancing.imbalance_percent • Introduced in: 9.11
rebalancing.runtime	string	query	False	Filter by rebalancing.runtime • Introduced in: 9.11
rebalancing.max_constituent_imbalance_percent	integer	query	False	Filter by rebalancing.max_constituent_imbalance_percent • Introduced in: 9.11
rebalancing.max_file_moves	integer	query	False	Filter by rebalancing.max_file_moves • Introduced in: 9.11
rebalancing.min_file_size	integer	query	False	Filter by rebalancing.min_file_size • Introduced in: 9.11

Name	Type	In	Required	Description
rebalancing.used_for_imbalance	integer	query	False	Filter by rebalancing.used_for_imbalance • Introduced in: 9.12
rebalancing.target_used	integer	query	False	Filter by rebalancing.target_used • Introduced in: 9.11
rebalancing.imbalance_size	integer	query	False	Filter by rebalancing.imbalance_size • Introduced in: 9.11
rebalancing.max_threshold	integer	query	False	Filter by rebalancing.max_threshold • Introduced in: 9.11
rebalancing.min_threshold	integer	query	False	Filter by rebalancing.min_threshold • Introduced in: 9.11
rebalancing.data_moved	integer	query	False	Filter by rebalancing.data_moved • Introduced in: 9.11
rebalancing.notices.arguments.code	string	query	False	Filter by rebalancing.notices.arguments.code • Introduced in: 9.12

Name	Type	In	Required	Description
rebalancing.notices.arguments.message	string	query	False	Filter by rebalancing.notices.arguments.message • Introduced in: 9.12
rebalancing.notices.message	string	query	False	Filter by rebalancing.notices.message • Introduced in: 9.12
rebalancing.notices.code	string	query	False	Filter by rebalancing.notices.code • Introduced in: 9.12
create_time	string	query	False	Filter by create_time
aggressive_readahead_mode	string	query	False	Filter by aggressive_readahead_mode • Introduced in: 9.13
status	string	query	False	Filter by status • Introduced in: 9.9
qos.policy.min_throughput_mbps	integer	query	False	Filter by qos.policy.min_throughput_mbps • Introduced in: 9.8 • Max value: 4194303 • Min value: 0

Name	Type	In	Required	Description
qos.policy.max_throughput_mbps	integer	query	False	Filter by qos.policy.max_throughput_mbps • Max value: 4194303 • Min value: 0
qos.policy.uuid	string	query	False	Filter by qos.policy.uuid
qos.policy.name	string	query	False	Filter by qos.policy.name
qos.policy.min_throughput_iops	integer	query	False	Filter by qos.policy.min_throughput_iops • Max value: 2147483647 • Min value: 0
qos.policy.max_throughput_iops	integer	query	False	Filter by qos.policy.max_throughput_iops • Max value: 2147483647 • Min value: 0
queue_for_encryption	boolean	query	False	Filter by queue_for_encryption • Introduced in: 9.8
space.cross_volume_dedupe_metafiles_footprint	integer	query	False	Filter by space.cross_volume_dedupe_metafiles_footprint • Introduced in: 9.10

Name	Type	In	Required	Description
space.volume_guarantee_footprint	integer	query	False	Filter by space.volume_guarantee_footprint • Introduced in: 9.10
space.physical_used	integer	query	False	Filter by space.physical_used • Introduced in: 9.10
space.max_size	string	query	False	Filter by space.max_size • Introduced in: 9.14
space.performance_tier_footprint	integer	query	False	Filter by space.performance_tier_footprint • Introduced in: 9.8
space.total_metadata	integer	query	False	Filter by space.total_metadata • Introduced in: 9.15
space.footprint	integer	query	False	Filter by space.footprint
space.capacity_tier_footprint	integer	query	False	Filter by space.capacity_tier_footprint
space.dedupe_metafiles_footprint	integer	query	False	Filter by space.dedupe_metafiles_footprint • Introduced in: 9.10

Name	Type	In	Required	Description
space.overwrite_reserve	integer	query	False	Filter by space.overwrite_reserve • Introduced in: 9.9
space.used	integer	query	False	Filter by space.used
space.fractional_reserve	integer	query	False	Filter by space.fractional_reserve • Introduced in: 9.9
space.capacity_tier_footprint_data_reduction	integer	query	False	Filter by space.capacity_tier_footprint_data_reduction • Introduced in: 9.13
space.auto_adaptive_compression_footprint_data_reduction	integer	query	False	Filter by space.auto_adaptive_compression_footprint_data_reduction • Introduced in: 9.11
space.logical_space.used_percent	integer	query	False	Filter by space.logical_space.used_percent • Introduced in: 9.9
space.logical_space.reporting	boolean	query	False	Filter by space.logical_space.reporting
space.logical_space.used_by_afs	integer	query	False	Filter by space.logical_space.used_by_afs

Name	Type	In	Required	Description
space.logical_space.enforcement	boolean	query	False	Filter by space.logical_space.enforcement
space.logical_space.available	integer	query	False	Filter by space.logical_space.available
space.logical_space.used	integer	query	False	Filter by space.logical_space.used • Introduced in: 9.9
space.logical_space.used_by_snapshots	integer	query	False	Filter by space.logical_space.used_by_snapshots • Introduced in: 9.10
space.large_size_enabled	boolean	query	False	Filter by space.large_size_enabled • Introduced in: 9.12
space.is_used_stale	boolean	query	False	Filter by space.is_used_stale • Introduced in: 9.12
space.snapshot.autodelete_trigger	string	query	False	Filter by space.snapshot.autodelete_trigger • Introduced in: 9.10
space.snapshot.reserve_size	integer	query	False	Filter by space.snapshot.reserve_size • Introduced in: 9.9

Name	Type	In	Required	Description
space.snapshot.space_used_percent	integer	query	False	Filter by space.snapshot.space_used_percent • Introduced in: 9.9
space.snapshot.autodelete.prefix	string	query	False	Filter by space.snapshot.autodelete.prefix • Introduced in: 9.13
space.snapshot.autodelete.trigger	string	query	False	Filter by space.snapshot.autodelete.trigger • Introduced in: 9.13
space.snapshot.autodelete.commitment	string	query	False	Filter by space.snapshot.autodelete.commitment • Introduced in: 9.13
space.snapshot.autodelete.delete_order	string	query	False	Filter by space.snapshot.autodelete.delete_order • Introduced in: 9.13
space.snapshot.autodelete.defer_delete	string	query	False	Filter by space.snapshot.autodelete.defer_delete • Introduced in: 9.13
space.snapshot.autodelete.target_free_space	integer	query	False	Filter by space.snapshot.autodelete.target_free_space • Introduced in: 9.13

Name	Type	In	Required	Description
space.snapshot.reserve_available	integer	query	False	Filter by space.snapshot.reserve_available • Introduced in: 9.10
space.snapshot.reserve_percent	integer	query	False	Filter by space.snapshot.reserve_percent
space.snapshot.used	integer	query	False	Filter by space.snapshot.used
space.metadata	integer	query	False	Filter by space.metadata
space.available	integer	query	False	Filter by space.available
space.file_operation_metadata	integer	query	False	Filter by space.file_operation_metadata • Introduced in: 9.10
space.physical_used_percent	integer	query	False	Filter by space.physical_used_percent • Introduced in: 9.10
space.afs_total	integer	query	False	Filter by space.afs_total • Introduced in: 9.9
space.snapmirror_destination_footprint	integer	query	False	Filter by space.snapmirror_destination_footprint • 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.compaction_footprint_data_reduction	integer	query	False	Filter by space.compaction_footprint_data_reduction • Introduced in: 9.13
space.local_tier_footprint	integer	query	False	Filter by space.local_tier_footprint • Introduced in: 9.8
space.over_provisioned	integer	query	False	Filter by space.over_provisioned
space.snapshot_spill	integer	query	False	Filter by space.snapshot_spill • Introduced in: 9.10
space.block_storage_inactive_user_data_percent	integer	query	False	Filter by space.block_storage_inactive_user_data_percent • Introduced in: 9.9
space.expected_available	integer	query	False	Filter by space.expected_available • Introduced in: 9.10

Name	Type	In	Required	Description
space.overwrite_reserve_used	integer	query	False	Filter by space.overwrite_reserve_used • Introduced in: 9.9
space.nearly_full_threshold_percent	integer	query	False	Filter by space.nearly_full_threshold_percent • Introduced in: 9.9
space.used_by_afs	integer	query	False	Filter by space.used_by_afs • Introduced in: 9.9
space.delayed_free_footprint	integer	query	False	Filter by space.delayed_free_footprint • Introduced in: 9.10
space.snapshot_reserve_unusable	integer	query	False	Filter by space.snapshot_reserve_unusable • Introduced in: 9.10
space.total_metadata_footprint	integer	query	False	Filter by space.total_metadata_footprint • Introduced in: 9.15
space.block_storage_inactive_user_data	integer	query	False	Filter by space.block_storage_inactive_user_data

Name	Type	In	Required	Description
space.cross_volume_dedupe_metafiles_temporary_footprint	integer	query	False	Filter by space.cross_volume_dedupe_metafiles_temporary_footprint • Introduced in: 9.10
space.dedupe_metafiles_temporary_footprint	integer	query	False	Filter by space.dedupe_metafiles_temporary_footprint • Introduced in: 9.10
space.total_footprint	integer	query	False	Filter by space.total_footprint • Introduced in: 9.8
space.user_data	integer	query	False	Filter by space.user_data • Introduced in: 9.10
space.percent_used	integer	query	False	Filter by space.percent_used • Introduced in: 9.9
space.size	integer	query	False	Filter by space.size
space.size_available_for_snapshots	integer	query	False	Filter by space.size_available_for_snapshots • Introduced in: 9.9
space.effective_total_footprint	integer	query	False	Filter by space.effective_total_footprint • Introduced in: 9.11

Name	Type	In	Required	Description
space.filesystem_size	integer	query	False	Filter by space.filesystem_size • Introduced in: 9.10
space.filesystem_size_fixed	boolean	query	False	Filter by space.filesystem_size_fixed • Introduced in: 9.10
space.available_percent	integer	query	False	Filter by space.available_percent • Introduced in: 9.9
msid	integer	query	False	Filter by msid • Introduced in: 9.11
type	string	query	False	Filter by type
svm.name	string	query	False	Filter by svm.name
svm.uuid	string	query	False	Filter by svm.uuid
autosize.maximum	integer	query	False	Filter by autosize.maximum
autosize.mode	string	query	False	Filter by autosize.mode
autosize.grow_threshold	integer	query	False	Filter by autosize.grow_threshold
autosize.shrink_threshold	integer	query	False	Filter by autosize.shrink_threshold

Name	Type	In	Required	Description
autosize.minimum	integer	query	False	Filter by autosize.minimum
clone.split_initiated	boolean	query	False	Filter by clone.split_initiated
clone.split_estimate	integer	query	False	Filter by clone.split_estimate
clone.is_flexclone	boolean	query	False	Filter by clone.is_flexclone
clone.has_flexclone	boolean	query	False	Filter by clone.has_flexclone • Introduced in: 9.16
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.parent_svm.name	string	query	False	Filter by clone.parent_svm.name
clone.parent_svm.uuid	string	query	False	Filter by clone.parent_svm.uuid
clone.split_complete_percent	integer	query	False	Filter by clone.split_complete_percent
clone.inherited_savings	integer	query	False	Filter by clone.inherited_savings • Introduced in: 9.12

Name	Type	In	Required	Description
clone.inherited_physical_used	integer	query	False	Filter by clone.inherited_physical_used • Introduced in: 9.12
clone.parent_snapshot.uuid	string	query	False	Filter by clone.parent_snapshot.uuid
clone.parent_snapshot.name	string	query	False	Filter by clone.parent_snapshot.name
nas.uid	integer	query	False	Filter by nas.uid
nas.unix_permissions	integer	query	False	Filter by nas.unix_permissions
nas.security_style	string	query	False	Filter by nas.security_style
nas.path	string	query	False	Filter by nas.path
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.export_policy.name	string	query	False	Filter by nas.export_policy.name
nas.export_policy.id	integer	query	False	Filter by nas.export_policy.id

Name	Type	In	Required	Description
nas.gid	integer	query	False	Filter by nas.gid
snaplock.compliance_clock_time	string	query	False	Filter by snaplock.compliance_clock_time
snaplock.autocommit_period	string	query	False	Filter by snaplock.autocommit_period
snaplock.expiry_time	string	query	False	Filter by snaplock.expiry_time
snaplock.is_audit_log	boolean	query	False	Filter by snaplock.is_audit_log
snaplock.append_mode_enabled	boolean	query	False	Filter by snaplock.append_mode_enabled
snaplock.retention.default	string	query	False	Filter by snaplock.retention.default
snaplock.retention.minimum	string	query	False	Filter by snaplock.retention.minimum
snaplock.retention.maximum	string	query	False	Filter by snaplock.retention.maximum
snaplock.type	string	query	False	Filter by snaplock.type
snaplock.litigation_count	integer	query	False	Filter by snaplock.litigation_count
snaplock.privileged_delete	string	query	False	Filter by snaplock.privileged_delete

Name	Type	In	Required	Description
snaplock.unspecified_retention_file_count	integer	query	False	Filter by snaplock.unspecified_retention_file_count • Introduced in: 9.8
efficiency.application_io_size	string	query	False	Filter by efficiency.application_io_size • Introduced in: 9.8
efficiency.progress	string	query	False	Filter by efficiency.progress • Introduced in: 9.9
efficiency.volume_path	string	query	False	Filter by efficiency.volume_path • Introduced in: 9.13
efficiency.last_op_err	string	query	False	Filter by efficiency.last_op_err • Introduced in: 9.9
efficiency.policy.name	string	query	False	Filter by efficiency.policy.name • Introduced in: 9.7
efficiency.state	string	query	False	Filter by efficiency.state • Introduced in: 9.9

Name	Type	In	Required	Description
efficiency.scanner.state	string	query	False	Filter by efficiency.scanner.state • Introduced in: 9.11
efficiency.scanner.compression	boolean	query	False	Filter by efficiency.scanner.compression • Introduced in: 9.11
efficiency.scanner.dedupe	boolean	query	False	Filter by efficiency.scanner.dedupe • Introduced in: 9.11
efficiency.scanner.scan_old_data	boolean	query	False	Filter by efficiency.scanner.scan_old_data • Introduced in: 9.11
efficiency.idcs_scanner.operation_state	string	query	False	Filter by efficiency.idcs_scanner.operation_state • Introduced in: 9.13
efficiency.idcs_scanner.mode	string	query	False	Filter by efficiency.idcs_scanner.mode • Introduced in: 9.13
efficiency.idcs_scanner.threshold_inactive_time	string	query	False	Filter by efficiency.idcs_scanner.threshold_inactive_time • Introduced in: 9.13

Name	Type	In	Required	Description
efficiency.idcs_scanner.enabled	boolean	query	False	Filter by efficiency.idcs_scanner.enabled • Introduced in: 9.13
efficiency.idcs_scanner.status	string	query	False	Filter by efficiency.idcs_scanner.status • Introduced in: 9.13
efficiency.auto_state	string	query	False	Filter by efficiency.auto_state • Introduced in: 9.12
efficiency.type	string	query	False	Filter by efficiency.type • Introduced in: 9.9
efficiency.cross_volume_dedupe	string	query	False	Filter by efficiency.cross_volume_dedupe
efficiency.last_op_size	integer	query	False	Filter by efficiency.last_op_size • Introduced in: 9.9
efficiency.has_savings	boolean	query	False	Filter by efficiency.has_savings • Introduced in: 9.11

Name	Type	In	Required	Description
efficiency.last_op_end	string	query	False	Filter by efficiency.last_op_end • Introduced in: 9.9
efficiency.storage_efficiency_mode	string	query	False	Filter by efficiency.storage_efficiency_mode • Introduced in: 9.10
efficiency.space_savings.dedupe_percent	integer	query	False	Filter by efficiency.space_savings.dedupe_percent • Introduced in: 9.11
efficiency.space_savings.total	integer	query	False	Filter by efficiency.space_savings.total • Introduced in: 9.11
efficiency.space_savings.dedupe	integer	query	False	Filter by efficiency.space_savings.dedupe • Introduced in: 9.11
efficiency.space_savings.compression	integer	query	False	Filter by efficiency.space_savings.compression • Introduced in: 9.11

Name	Type	In	Required	Description
efficiency.space_savings.compression_percent	integer	query	False	Filter by efficiency.space_savings.compression_percent • Introduced in: 9.11
efficiency.space_savings.total_percent	integer	query	False	Filter by efficiency.space_savings.total_percent • Introduced in: 9.11
efficiency.space_savings.dedupe_sharing	integer	query	False	Filter by efficiency.space_savings.dedupe_sharing • Introduced in: 9.11
efficiency.last_op_begin	string	query	False	Filter by efficiency.last_op_begin • Introduced in: 9.9
efficiency.compression	string	query	False	Filter by efficiency.compression
efficiency.logging_enabled	boolean	query	False	Filter by efficiency.logging_enabled • Introduced in: 9.11
efficiency.op_state	string	query	False	Filter by efficiency.op_state • Introduced in: 9.9

Name	Type	In	Required	Description
efficiency.last_op_state	string	query	False	Filter by efficiency.last_op_state • Introduced in: 9.9
efficiency.schedule	string	query	False	Filter by efficiency.schedule • Introduced in: 9.8
efficiency.compression_type	string	query	False	Filter by efficiency.compression_type • Introduced in: 9.11
efficiency.dedupe	string	query	False	Filter by efficiency.dedupe
efficiency.compaction	string	query	False	Filter by efficiency.compaction
asynchronous_directory_delete.trash_bin	string	query	False	Filter by asynchronous_directory_delete.trash_bin • Introduced in: 9.11
asynchronous_directory_delete.enabled	boolean	query	False	Filter by asynchronous_directory_delete.enabled • Introduced in: 9.11
movement.destination_aggregate.uuid	string	query	False	Filter by movement.destination_aggregate.uuid
movement.destination_aggregate.name	string	query	False	Filter by movement.destination_aggregate.name

Name	Type	In	Required	Description
movement.percent_complete	integer	query	False	Filter by movement.percent_complete
movement.cutover_window	integer	query	False	Filter by movement.cutover_window
movement.state	string	query	False	Filter by movement.state
movement.start_time	string	query	False	Filter by movement.start_time • Introduced in: 9.9
_tags	string	query	False	Filter by _tags • Introduced in: 9.13
flexcache_endpoint_type	string	query	False	Filter by flexcache_endpoint_type
size	integer	query	False	Filter by size
snapmirror.is_protected	boolean	query	False	Filter by snapmirror.is_protected • Introduced in: 9.7
snapmirror.destinations.is_ontap	boolean	query	False	Filter by snapmirror.destinations.is_ontap • Introduced in: 9.9

Name	Type	In	Required	Description
snapmirror.destinations.is_cloud	boolean	query	False	Filter by snapmirror.destinations.is_cloud • Introduced in: 9.9
comment	string	query	False	Filter by comment • maxLength: 1023 • minLength: 0
quota.state	string	query	False	Filter by quota.state
application.uuid	string	query	False	Filter by application.uuid
application.name	string	query	False	Filter by application.name
cloud_retrieval_policy	string	query	False	Filter by cloud_retrieval_policy • Introduced in: 9.8
tiering.policy	string	query	False	Filter by tiering.policy
tiering.object_tags	string	query	False	Filter by tiering.object_tags • Introduced in: 9.8 • maxLength: 257
tiering.min_cooling_days	integer	query	False	Filter by tiering.min_cooling_days • Introduced in: 9.8 • Max value: 183 • Min value: 2

Name	Type	In	Required	Description
statistics.timestamp	string	query	False	Filter by statistics.timestamp
statistics.latency_raw.write	integer	query	False	Filter by statistics.latency_raw.write
statistics.latency_raw.total	integer	query	False	Filter by statistics.latency_raw.total
statistics.latency_raw.read	integer	query	False	Filter by statistics.latency_raw.read
statistics.latency_raw.other	integer	query	False	Filter by statistics.latency_raw.other
statistics.status	string	query	False	Filter by statistics.status
statistics.iops_raw.write	integer	query	False	Filter by statistics.iops_raw.write
statistics.iops_raw.total	integer	query	False	Filter by statistics.iops_raw.total
statistics.iops_raw.read	integer	query	False	Filter by statistics.iops_raw.read
statistics.iops_raw.other	integer	query	False	Filter by statistics.iops_raw.other
statistics.nfs_ops_raw.watch.count	integer	query	False	Filter by statistics.nfs_ops_raw.watch.count • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.nfs_ops_raw.watch.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.watch.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.readdir.count	integer	query	False	Filter by statistics.nfs_ops_raw.readdir.count • Introduced in: 9.11
statistics.nfs_ops_raw.readdir.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.readdir.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.setattr.count	integer	query	False	Filter by statistics.nfs_ops_raw.setattr.count • Introduced in: 9.11
statistics.nfs_ops_raw.setattr.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.setattr.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.write.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.write.total_time • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.nfs_ops_raw.write.volume_protocol_size_histogram_counts	integer	query	False	Filter by statistics.nfs_ops_raw.write.volume_protocol_size_histogram_counts • Introduced in: 9.11
statistics.nfs_ops_raw.write.volume_protocol_latency_histogram_counts	integer	query	False	Filter by statistics.nfs_ops_raw.write.volume_protocol_latency_histogram_counts • Introduced in: 9.11
statistics.nfs_ops_raw.write.count	integer	query	False	Filter by statistics.nfs_ops_raw.write.count • Introduced in: 9.11
statistics.nfs_ops_raw.write.volume_protocol_size_histogram_labels	string	query	False	Filter by statistics.nfs_ops_raw.write.volume_protocol_size_histogram_labels • Introduced in: 9.11
statistics.nfs_ops_raw.write.volume_protocol_latency_histogram_labels	string	query	False	Filter by statistics.nfs_ops_raw.write.volume_protocol_latency_histogram_labels • Introduced in: 9.11
statistics.nfs_ops_raw.create.symlink.count	integer	query	False	Filter by statistics.nfs_ops_raw.create.symlink.count • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.nfs_ops_raw.create.symlink.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.create.symlink.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.create.file.count	integer	query	False	Filter by statistics.nfs_ops_raw.create.file.count • Introduced in: 9.11
statistics.nfs_ops_raw.create.file.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.create.file.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.create.other.count	integer	query	False	Filter by statistics.nfs_ops_raw.create.other.count • Introduced in: 9.11
statistics.nfs_ops_raw.create.other.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.create.other.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.create.dir.count	integer	query	False	Filter by statistics.nfs_ops_raw.create.dir.count • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.nfs_ops_raw.create.dir.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.create.dir.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.getattr.count	integer	query	False	Filter by statistics.nfs_ops_raw.getattr.count • Introduced in: 9.11
statistics.nfs_ops_raw.getattr.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.getattr.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.lock.count	integer	query	False	Filter by statistics.nfs_ops_raw.lock.count • Introduced in: 9.11
statistics.nfs_ops_raw.lock.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.lock.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.audit.count	integer	query	False	Filter by statistics.nfs_ops_raw.audit.count • Introduced in: 9.11
statistics.nfs_ops_raw.audit.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.audit.total_time • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.nfs_ops_raw.link.count	integer	query	False	Filter by statistics.nfs_ops_raw.link.count • Introduced in: 9.11
statistics.nfs_ops_raw.link.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.link.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.readlink.count	integer	query	False	Filter by statistics.nfs_ops_raw.readlink.count • Introduced in: 9.11
statistics.nfs_ops_raw.readlink.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.readlink.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.read.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.read.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.read.volume_protocol_size_histogram_counts	integer	query	False	Filter by statistics.nfs_ops_raw.read.volume_protocol_size_histogram_counts • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.nfs_ops_raw.read.volume_protocol_latency_histogram_counts	integer	query	False	Filter by statistics.nfs_ops_raw.read.volume_protocol_latency_histogram_counts • Introduced in: 9.11
statistics.nfs_ops_raw.read.count	integer	query	False	Filter by statistics.nfs_ops_raw.read.count • Introduced in: 9.11
statistics.nfs_ops_raw.read.volume_protocol_size_histogram_labels	string	query	False	Filter by statistics.nfs_ops_raw.read.volume_protocol_size_histogram_labels • Introduced in: 9.11
statistics.nfs_ops_raw.read.volume_protocol_latency_histogram_labels	string	query	False	Filter by statistics.nfs_ops_raw.read.volume_protocol_latency_histogram_labels • Introduced in: 9.11
statistics.nfs_ops_raw.open.count	integer	query	False	Filter by statistics.nfs_ops_raw.open.count • Introduced in: 9.11
statistics.nfs_ops_raw.open.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.open.total_time • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.nfs_ops_raw.access.count	integer	query	False	Filter by statistics.nfs_ops_raw.access.count • Introduced in: 9.11
statistics.nfs_ops_raw.access.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.access.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.unlink.count	integer	query	False	Filter by statistics.nfs_ops_raw.unlink.count • Introduced in: 9.11
statistics.nfs_ops_raw.unlink.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.unlink.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.lookup.count	integer	query	False	Filter by statistics.nfs_ops_raw.lookup.count • Introduced in: 9.11
statistics.nfs_ops_raw.lookup.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.lookup.total_time • Introduced in: 9.11
statistics.nfs_ops_raw.rename.count	integer	query	False	Filter by statistics.nfs_ops_raw.rename.count • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.nfs_ops_raw.rename.total_time	integer	query	False	Filter by statistics.nfs_ops_raw.rename.total_time • Introduced in: 9.11
statistics.flexcache_raw.cache_miss_blocks	integer	query	False	Filter by statistics.flexcache_raw.cache_miss_blocks • Introduced in: 9.8
statistics.flexcache_raw.client_requested_blocks	integer	query	False	Filter by statistics.flexcache_raw.client_requested_blocks • Introduced in: 9.8
statistics.flexcache_raw.status	string	query	False	Filter by statistics.flexcache_raw.status • Introduced in: 9.8
statistics.flexcache_raw.timestamp	string	query	False	Filter by statistics.flexcache_raw.timestamp • Introduced in: 9.8
statistics.throughput_raw.write	integer	query	False	Filter by statistics.throughput_raw.write
statistics.throughput_raw.total	integer	query	False	Filter by statistics.throughput_raw.total
statistics.throughput_raw.read	integer	query	False	Filter by statistics.throughput_raw.read

Name	Type	In	Required	Description
statistics.throughput_raw.other	integer	query	False	Filter by statistics.throughput_raw.other
statistics.cloud.timestamp	string	query	False	Filter by statistics.cloud.timestamp • 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.total	integer	query	False	Filter by statistics.cloud.latency_raw.total • Introduced in: 9.7
statistics.cloud.latency_raw.read	integer	query	False	Filter by statistics.cloud.latency_raw.read • 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.status	string	query	False	Filter by statistics.cloud.status • Introduced in: 9.7

Name	Type	In	Required	Description
statistics.cloud.iops_raw.write	integer	query	False	Filter by statistics.cloud.iops_raw.write • Introduced in: 9.7
statistics.cloud.iops_raw.total	integer	query	False	Filter by statistics.cloud.iops_raw.total • Introduced in: 9.7
statistics.cloud.iops_raw.read	integer	query	False	Filter by statistics.cloud.iops_raw.read • Introduced in: 9.7
statistics.cloud.iops_raw.other	integer	query	False	Filter by statistics.cloud.iops_raw.other • Introduced in: 9.7
statistics.cifs_ops_raw.watch.count	integer	query	False	Filter by statistics.cifs_ops_raw.watch.count • Introduced in: 9.11
statistics.cifs_ops_raw.watch.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.watch.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.readdir.count	integer	query	False	Filter by statistics.cifs_ops_raw.readdir.count • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.cifs_ops_raw.readdir.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.readdir.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.setattr.count	integer	query	False	Filter by statistics.cifs_ops_raw.setattr.count • Introduced in: 9.11
statistics.cifs_ops_raw.setattr.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.setattr.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.write.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.write.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.write.volume_protocol_size_histogram_counts	integer	query	False	Filter by statistics.cifs_ops_raw.write.volume_protocol_size_histogram_counts • Introduced in: 9.11
statistics.cifs_ops_raw.write.volume_protocol_latency_histogram_counts	integer	query	False	Filter by statistics.cifs_ops_raw.write.volume_protocol_latency_histogram_counts • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.cifs_ops_raw.write.count	integer	query	False	Filter by statistics.cifs_ops_raw.write.count • Introduced in: 9.11
statistics.cifs_ops_raw.write.volume_protocol_size_histogram_labels	string	query	False	Filter by statistics.cifs_ops_raw.write.volume_protocol_size_histogram_labels • Introduced in: 9.11
statistics.cifs_ops_raw.write.volume_protocol_latency_histogram_labels	string	query	False	Filter by statistics.cifs_ops_raw.write.volume_protocol_latency_histogram_labels • Introduced in: 9.11
statistics.cifs_ops_raw.create.symlink.count	integer	query	False	Filter by statistics.cifs_ops_raw.create.symlink.count • Introduced in: 9.11
statistics.cifs_ops_raw.create.symlink.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.create.symlink.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.create.file.count	integer	query	False	Filter by statistics.cifs_ops_raw.create.file.count • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.cifs_ops_raw.create.file.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.create.file.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.create.other.count	integer	query	False	Filter by statistics.cifs_ops_raw.create.other.count • Introduced in: 9.11
statistics.cifs_ops_raw.create.other.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.create.other.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.create.dir.count	integer	query	False	Filter by statistics.cifs_ops_raw.create.dir.count • Introduced in: 9.11
statistics.cifs_ops_raw.create.dir.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.create.dir.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.getattr.count	integer	query	False	Filter by statistics.cifs_ops_raw.getattr.count • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.cifs_ops_raw.getattr.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.getattr.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.lock.count	integer	query	False	Filter by statistics.cifs_ops_raw.lock.count • Introduced in: 9.11
statistics.cifs_ops_raw.lock.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.lock.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.audit.count	integer	query	False	Filter by statistics.cifs_ops_raw.audit.count • Introduced in: 9.11
statistics.cifs_ops_raw.audit.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.audit.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.link.count	integer	query	False	Filter by statistics.cifs_ops_raw.link.count • Introduced in: 9.11
statistics.cifs_ops_raw.link.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.link.total_time • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.cifs_ops_raw.readlink.count	integer	query	False	Filter by statistics.cifs_ops_raw.readlink.count • Introduced in: 9.11
statistics.cifs_ops_raw.readlink.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.readlink.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.read.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.read.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.read.volume_protocol_size_histogram_counts	integer	query	False	Filter by statistics.cifs_ops_raw.read.volume_protocol_size_histogram_counts • Introduced in: 9.11
statistics.cifs_ops_raw.read.volume_protocol_latency_histogram_counts	integer	query	False	Filter by statistics.cifs_ops_raw.read.volume_protocol_latency_histogram_counts • Introduced in: 9.11
statistics.cifs_ops_raw.read.count	integer	query	False	Filter by statistics.cifs_ops_raw.read.count • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.cifs_ops_raw.read.volume_protocol_size_histogram_labels	string	query	False	Filter by statistics.cifs_ops_raw.read.volume_protocol_size_histogram_labels • Introduced in: 9.11
statistics.cifs_ops_raw.read.volume_protocol_latency_histogram_labels	string	query	False	Filter by statistics.cifs_ops_raw.read.volume_protocol_latency_histogram_labels • Introduced in: 9.11
statistics.cifs_ops_raw.open.count	integer	query	False	Filter by statistics.cifs_ops_raw.open.count • Introduced in: 9.11
statistics.cifs_ops_raw.open.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.open.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.access.count	integer	query	False	Filter by statistics.cifs_ops_raw.access.count • Introduced in: 9.11
statistics.cifs_ops_raw.access.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.access.total_time • Introduced in: 9.11

Name	Type	In	Required	Description
statistics.cifs_ops_raw.unlink.count	integer	query	False	Filter by statistics.cifs_ops_raw.unlink.count • Introduced in: 9.11
statistics.cifs_ops_raw.unlink.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.unlink.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.lookup.count	integer	query	False	Filter by statistics.cifs_ops_raw.lookup.count • Introduced in: 9.11
statistics.cifs_ops_raw.lookup.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.lookup.total_time • Introduced in: 9.11
statistics.cifs_ops_raw.rename.count	integer	query	False	Filter by statistics.cifs_ops_raw.rename.count • Introduced in: 9.11
statistics.cifs_ops_raw.rename.total_time	integer	query	False	Filter by statistics.cifs_ops_raw.rename.total_time • Introduced in: 9.11
snapshot_count	integer	query	False	Filter by snapshot_count • Introduced in: 9.10 • Max value: 1023 • Min value: 0

Name	Type	In	Required	Description
state	string	query	False	Filter by state
guarantee.type	string	query	False	Filter by guarantee.type
guarantee.honored	boolean	query	False	Filter by guarantee.honored
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
is_object_store	boolean	query	False	Filter by is_object_store • Introduced in: 9.8
encryption.key_create_time	string	query	False	Filter by encryption.key_create_time • Introduced in: 9.11
encryption.type	string	query	False	Filter by encryption.type
encryption.status.code	string	query	False	Filter by encryption.status.code

Name	Type	In	Required	Description
encryption.status.message	string	query	False	Filter by encryption.status.message
encryption.enabled	boolean	query	False	Filter by encryption.enabled
encryption.state	string	query	False	Filter by encryption.state
encryption.action	string	query	False	Filter by encryption.action • Introduced in: 9.14
encryption.rekey	boolean	query	False	Filter by encryption.rekey
encryption.key_id	string	query	False	Filter by encryption.key_id
uuid	string	query	False	Filter by uuid
snapshot_locking_enabled	boolean	query	False	Filter by snapshot_locking_enabled • Introduced in: 9.12
activity_tracking.supported	boolean	query	False	Filter by activity_tracking.supported • Introduced in: 9.10
activity_tracking.unsupported_reason.message	string	query	False	Filter by activity_tracking.unsupported_reason.message • Introduced in: 9.10

Name	Type	In	Required	Description
activity_tracking.unsupported_reason.code	string	query	False	Filter by activity_tracking.unsupported_reason.code • Introduced in: 9.10
activity_tracking.state	string	query	False	Filter by activity_tracking.state • Introduced in: 9.10
granular_data	boolean	query	False	Filter by granular_data • Introduced in: 9.11
flexgroup.name	string	query	False	Filter by flexgroup.name • Introduced in: 9.10 • maxLength: 203 • minLength: 1
flexgroup.uuid	string	query	False	Filter by flexgroup.uuid • Introduced in: 9.10
metric.flexcache.status	string	query	False	Filter by metric.flexcache.status • Introduced in: 9.8
metric.flexcache.cache_miss_percent	integer	query	False	Filter by metric.flexcache.cache_miss_percent • Introduced in: 9.8

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

Name	Type	In	Required	Description
metric.iops.read	integer	query	False	Filter by metric.iops.read
metric.iops.other	integer	query	False	Filter by metric.iops.other
metric.latency.write	integer	query	False	Filter by metric.latency.write
metric.latency.total	integer	query	False	Filter by metric.latency.total
metric.latency.read	integer	query	False	Filter by metric.latency.read
metric.latency.other	integer	query	False	Filter by metric.latency.other
metric.cloud.latency.write	integer	query	False	Filter by metric.cloud.latency.write • Introduced in: 9.7
metric.cloud.latency.total	integer	query	False	Filter by metric.cloud.latency.total • Introduced in: 9.7
metric.cloud.latency.read	integer	query	False	Filter by metric.cloud.latency.read • Introduced in: 9.7
metric.cloud.latency.other	integer	query	False	Filter by metric.cloud.latency.other • Introduced in: 9.7

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

Name	Type	In	Required	Description
metric.duration	string	query	False	Filter by metric.duration
name	string	query	False	Filter by name • maxLength: 203 • minLength: 1
access_time_enabled	boolean	query	False	Filter by access_time_enabled • Introduced in: 9.8
consistency_group.uuid	string	query	False	Filter by consistency_group.uuid • Introduced in: 9.10
consistency_group.name	string	query	False	Filter by consistency_group.name • Introduced in: 9.7
snapshot_directory_access_enabled	boolean	query	False	Filter by snapshot_directory_access_enabled • Introduced in: 9.13
validate_only	boolean	query	False	Filter by validate_only • Introduced in: 9.14
is_svm_root	boolean	query	False	Filter by is_svm_root • Introduced in: 9.7

Name	Type	In	Required	Description
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
analytics.initialization.state	string	query	False	Filter by analytics.initialization.state • Introduced in: 9.12
analytics.supported	boolean	query	False	Filter by analytics.supported • Introduced in: 9.8
analytics.total_files	integer	query	False	Filter by analytics.total_files • Introduced in: 9.14
analytics.files_scanned	integer	query	False	Filter by analytics.files_scanned • Introduced in: 9.14
analytics.state	string	query	False	Filter by analytics.state • Introduced in: 9.8

Name	Type	In	Required	Description
analytics.scan_throttle_reason.code	string	query	False	Filter by analytics.scan_throttle_reason.code • Introduced in: 9.14
analytics.scan_throttle_reason.message	string	query	False	Filter by analytics.scan_throttle_reason.message • Introduced in: 9.14
analytics.scan_throttle_reason.arguments	string	query	False	Filter by analytics.scan_throttle_reason.arguments • Introduced in: 9.14
analytics.scan_progress	integer	query	False	Filter by analytics.scan_progress • Introduced in: 9.8
files.maximum	integer	query	False	Filter by files.maximum
files.used	integer	query	False	Filter by files.used
granular_data_mode	string	query	False	Filter by granular_data_mode • Introduced in: 9.16
max_dir_size	integer	query	False	Filter by max_dir_size • Introduced in: 9.10
snapshot_policy.uuid	string	query	False	Filter by snapshot_policy.uuid

Name	Type	In	Required	Description
snapshot_policy.name	string	query	False	Filter by snapshot_policy.name
anti_ransomware.workload.historical_statistics.high_entropy_data_write_peak_percent	integer	query	False	Filter by anti_ransomware.workload.historical_statistics.high_entropy_data_write_peak_percent • Introduced in: 9.16
anti_ransomware.workload.historical_statistics.high_entropy_data_write_peak_rate_kb_per_minute	integer	query	False	Filter by anti_ransomware.workload.historical_statistics.high_entropy_data_write_peak_rate_kb_per_minute • Introduced in: 9.16
anti_ransomware.workload.historical_statistics.file_delete_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.workload.historical_statistics.file_delete_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.workload.historical_statistics.file_rename_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.workload.historical_statistics.file_rename_peak_rate_per_minute • Introduced in: 9.16

Name	Type	In	Required	Description
anti_ransomware.workload.historical_statistics.file_create_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.workload.historical_statistics.file_create_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.workload.surge_statistics.time	string	query	False	Filter by anti_ransomware.workload.surge_statistics.time • Introduced in: 9.16
anti_ransomware.workload.surge_statistics.file_create_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.workload.surge_statistics.file_create_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.workload.surge_statistics.high_entropy_data_write_peak_percent	integer	query	False	Filter by anti_ransomware.workload.surge_statistics.high_entropy_data_write_peak_percent • Introduced in: 9.16
anti_ransomware.workload.surge_statistics.high_entropy_data_write_peak_rate_kb_per_minute	integer	query	False	Filter by anti_ransomware.workload.surge_statistics.high_entropy_data_write_peak_rate_kb_per_minute • Introduced in: 9.16

Name	Type	In	Required	Description
anti_ransomware.workload.surge_statistics.file_rename_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.workload.surge_statistics.file_rename_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.workload.surge_statistics.file_delete_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.workload.surge_statistics.file_delete_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.workload.newly_observed_file_extensions.count	integer	query	False	Filter by anti_ransomware.workload.newly_observed_file_extensions.count • Introduced in: 9.16
anti_ransomware.workload.newly_observed_file_extensions.name	string	query	False	Filter by anti_ransomware.workload.newly_observed_file_extensions.name • Introduced in: 9.16
anti_ransomware.workload.typical_usage.file_rename_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.workload.typical_usage.file_rename_peak_rate_per_minute • Introduced in: 9.16

Name	Type	In	Required	Description
anti_ransomware.workload.typical_usage.file_delete_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.workload.typical_usage.file_delete_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.workload.typical_usage.high_entropy_data_write_peak_rate_kb_per_minute	integer	query	False	Filter by anti_ransomware.workload.typical_usage.high_entropy_data_write_peak_rate_kb_per_minute • Introduced in: 9.16
anti_ransomware.workload.typical_usage.high_entropy_data_write_peak_percent	integer	query	False	Filter by anti_ransomware.workload.typical_usage.high_entropy_data_write_peak_percent • Introduced in: 9.16
anti_ransomware.workload.typical_usage.file_create_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.workload.typical_usage.file_create_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.workload.file_extensions_observed	string	query	False	Filter by anti_ransomware.workload.file_extensions_observed • Introduced in: 9.16

Name	Type	In	Required	Description
anti_ransomware.workload.file_extension_types_count	integer	query	False	Filter by anti_ransomware.workload.file_extension_types_count • Introduced in: 9.16
anti_ransomware.workload.surge_usage.file_rename_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.workload.surge_usage.file_rename_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.workload.surge_usage.file_delete_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.workload.surge_usage.file_delete_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.workload.surge_usage.high_entropy_data_write_peak_percent	integer	query	False	Filter by anti_ransomware.workload.surge_usage.high_entropy_data_write_peak_percent • Introduced in: 9.16
anti_ransomware.workload.surge_usage.high_entropy_data_write_peak_rate_kb_per_minute	integer	query	False	Filter by anti_ransomware.workload.surge_usage.high_entropy_data_write_peak_rate_kb_per_minute • Introduced in: 9.16

Name	Type	In	Required	Description
anti_ransomware.workload.surge_usage.file_create_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.workload.surge_usage.file_create_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.workload.surge_usage.time	string	query	False	Filter by anti_ransomware.workload.surge_usage.time • Introduced in: 9.16
anti_ransomware.workload.surge_usage.newly_observed_file_extensions.name	string	query	False	Filter by anti_ransomware.workload.surge_usage.newly_observed_file_extensions.name • Introduced in: 9.16
anti_ransomware.workload.surge_usage.newly_observed_file_extensions.count	integer	query	False	Filter by anti_ransomware.workload.surge_usage.newly_observed_file_extensions.count • Introduced in: 9.16
anti_ransomware.surge_usage.file_create_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.surge_usage.file_create_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.surge_usage.time	string	query	False	Filter by anti_ransomware.surge_usage.time • Introduced in: 9.16

Name	Type	In	Required	Description
anti_ransomware.surge_usage.file_rename_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.surge_usage.file_rename_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.surge_usage.file_delete_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.surge_usage.file_delete_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.surge_usage.high_entropy_data_write_peak_rate_kb_per_minute	integer	query	False	Filter by anti_ransomware.surge_usage.high_entropy_data_write_peak_rate_kb_per_minute • Introduced in: 9.16
anti_ransomware.surge_usage.high_entropy_data_write_peak_percent	integer	query	False	Filter by anti_ransomware.surge_usage.high_entropy_data_write_peak_percent • Introduced in: 9.16
anti_ransomware.update_baseline_from_surge	boolean	query	False	Filter by anti_ransomware.update_baseline_from_surge • Introduced in: 9.15

Name	Type	In	Required	Description
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.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.surge_as_normal	boolean	query	False	Filter by anti_ransomware.surge_as_normal • Introduced in: 9.11
anti_ransomware.suspect_files.format	string	query	False	Filter by anti_ransomware.suspect_files.format • Introduced in: 9.10
anti_ransomware.suspect_files.count	integer	query	False	Filter by anti_ransomware.suspect_files.count • Introduced in: 9.10

Name	Type	In	Required	Description
anti_ransomware.suspect_files.entropy	string	query	False	Filter by anti_ransomware.suspect_files.entropy • Introduced in: 9.11
anti_ransomware.state	string	query	False	Filter by anti_ransomware.state • Introduced in: 9.10
anti_ransomware.attack_reports.time	string	query	False	Filter by anti_ransomware.attack_reports.time • Introduced in: 9.10
anti_ransomware.attack_probability	string	query	False	Filter by anti_ransomware.attack_probability • Introduced in: 9.10
anti_ransomware.attack_detection_parameters.file_rename_op_rate_surge_notify_percent	integer	query	False	Filter by anti_ransomware.attack_detection_parameters.file_rename_op_rate_surge_notify_percent • Introduced in: 9.16
anti_ransomware.attack_detection_parameters.high_entropy_data_surge_notify_percent	integer	query	False	Filter by anti_ransomware.attack_detection_parameters.high_entropy_data_surge_notify_percent • Introduced in: 9.16

Name	Type	In	Required	Description
anti_ransomware.attack_detection_parameters.never_seen_before_file_extension_duration_in_hours	integer	query	False	Filter by anti_ransomware.attack_detection_parameters.never_seen_before_file_extension_duration_in_hours • Introduced in: 9.16
anti_ransomware.attack_detection_parameters.based_on_never_seen_before_file_extension	boolean	query	False	Filter by anti_ransomware.attack_detection_parameters.based_on_never_seen_before_file_extension • Introduced in: 9.16
anti_ransomware.attack_detection_parameters.based_on_high_entropy_data_rate	boolean	query	False	Filter by anti_ransomware.attack_detection_parameters.based_on_high_entropy_data_rate • Introduced in: 9.16
anti_ransomware.attack_detection_parameters.never_seen_before_file_extension_count_notify_threshold	integer	query	False	Filter by anti_ransomware.attack_detection_parameters.never_seen_before_file_extension_count_notify_threshold • Introduced in: 9.16
anti_ransomware.attack_detection_parameters.file_delete_op_rate_surge_notify_percent	integer	query	False	Filter by anti_ransomware.attack_detection_parameters.file_delete_op_rate_surge_notify_percent • Introduced in: 9.16

Name	Type	In	Required	Description
anti_ransomware.attack_detection_parameters.based_on_file_create_op_rate	boolean	query	False	Filter by anti_ransomware.attack_detection_parameters.based_on_file_create_op_rate • Introduced in: 9.16
anti_ransomware.attack_detection_parameters.relaxing_popular_file_extensions	boolean	query	False	Filter by anti_ransomware.attack_detection_parameters.relaxing_popular_file_extensions • Introduced in: 9.16
anti_ransomware.attack_detection_parameters.file_create_op_rate_surge_notify_percent	integer	query	False	Filter by anti_ransomware.attack_detection_parameters.file_create_op_rate_surge_notify_percent • Introduced in: 9.16
anti_ransomware.attack_detection_parameters.based_on_file_delete_op_rate	boolean	query	False	Filter by anti_ransomware.attack_detection_parameters.based_on_file_delete_op_rate • Introduced in: 9.16
anti_ransomware.attack_detection_parameters.based_on_file_rename_op_rate	boolean	query	False	Filter by anti_ransomware.attack_detection_parameters.based_on_file_rename_op_rate • Introduced in: 9.16

Name	Type	In	Required	Description
anti_ransomware.dry_run_start_time	string	query	False	Filter by anti_ransomware.dry_run_start_time • Introduced in: 9.10
anti_ransomware.event_log.is_enabled_on_snapshot_copy_creation	boolean	query	False	Filter by anti_ransomware.event_log.is_enabled_on_snapshot_copy_creation • Introduced in: 9.14
anti_ransomware.event_log.is_enabled_on_new_file_extension_seen	boolean	query	False	Filter by anti_ransomware.event_log.is_enabled_on_new_file_extension_seen • Introduced in: 9.14
anti_ransomware.typical_usage.file_create_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.typical_usage.file_create_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.typical_usage.file_delete_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.typical_usage.file_delete_peak_rate_per_minute • Introduced in: 9.16

Name	Type	In	Required	Description
anti_ransomware.typical_usage.file_rename_peak_rate_per_minute	integer	query	False	Filter by anti_ransomware.typical_usage.file_rename_peak_rate_per_minute • Introduced in: 9.16
anti_ransomware.typical_usage.high_entropy_data_write_peak_percent	integer	query	False	Filter by anti_ransomware.typical_usage.high_entropy_data_write_peak_percent • Introduced in: 9.16
anti_ransomware.typical_usage.high_entropy_data_write_peak_rate_kb_per_minute	integer	query	False	Filter by anti_ransomware.typical_usage.high_entropy_data_write_peak_rate_kb_per_minute • Introduced in: 9.16
aggregates.uuid	string	query	False	Filter by aggregates.uuid
aggregates.name	string	query	False	Filter by aggregates.name
scheduled_snapshot_naming_scheme	string	query	False	Filter by scheduled_snapshot_naming_scheme • Introduced in: 9.10
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

```
Status: 200, Ok
```

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

Example response


```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "_tags": [
        "team:csi",
        "environment:test"
      ],
      "activity_tracking": {
        "state": "string",
        "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"
        }
      ],
      "aggressive_readahead_mode": "string",
      "analytics": {
        "files_scanned": 43002,
        "initialization": {
```

```
          "state": "string"
        },
        "scan_progress": 17,
        "scan_throttle_reason": {
          "arguments": [
            "string"
          ],
          "code": "6739881",
          "message": "The file system analytics scan running on volume
\"fsavol2\" in SVM \"vs2\" has slowed down on node \"my_node\". Reason:
Computing resources are being used by higher priority workloads."
        },
        "state": "string",
        "total_files": 101890,
        "unsupported_reason": {
          "code": "111411207",
          "message": "File system analytics cannot be enabled on
volumes that contain LUNs."
        }
      },
      "anti_ransomware": {
        "attack_detection_parameters": {
          "file_create_op_rate_surge_notify_percent": 100
        },
        "attack_probability": "string",
        "attack_reports": [
          {
            "_links": {
              "suspects": {
                "href": "/api/resourcelink"
              }
            },
            "time": "2021-06-01 11:06:41 -0400"
          }
        ],
        "dry_run_start_time": "string",
        "space": {
          "snapshot_count": 0,
          "used": 0,
          "used_by_logs": 0,
          "used_by_snapshots": 0
        },
        "state": "string",
        "surge_usage": {
          "file_create_peak_rate_per_minute": 10,
          "file_delete_peak_rate_per_minute": 50,
```

```
        "file_rename_peak_rate_per_minute": 30,
        "high_entropy_data_write_peak_percent": 30,
        "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
        "time": "2021-12-01 12:46:20 -0500"
      },
      "suspect_files": [
        {
          "count": 0,
          "entropy": "string",
          "format": "string"
        }
      ],
      "typical_usage": {
        "file_create_peak_rate_per_minute": 50,
        "file_delete_peak_rate_per_minute": 10,
        "file_rename_peak_rate_per_minute": 5,
        "high_entropy_data_write_peak_percent": 10,
        "high_entropy_data_write_peak_rate_kb_per_minute": 1200
      },
      "workload": {
        "file_extension_types_count": 3,
        "file_extensions_observed": [
          "pdf",
          "jpeg",
          "txt"
        ],
        "historical_statistics": {
          "file_create_peak_rate_per_minute": 50,
          "file_delete_peak_rate_per_minute": 10,
          "file_rename_peak_rate_per_minute": 5,
          "high_entropy_data_write_peak_percent": 10,
          "high_entropy_data_write_peak_rate_kb_per_minute": 1200
        },
        "newly_observed_file_extensions": [
          {
            "count": [
              20
            ],
            "name": [
              "lockile"
            ]
          }
        ],
        "surge_statistics": {
          "file_create_peak_rate_per_minute": 10,
          "file_delete_peak_rate_per_minute": 50,
```

```
          "file_rename_peak_rate_per_minute": 30,
          "high_entropy_data_write_peak_percent": 30,
          "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
          "time": "2021-12-01 12:46:20 -0500"
        },
        "surge_usage": {
          "file_create_peak_rate_per_minute": 10,
          "file_delete_peak_rate_per_minute": 50,
          "file_rename_peak_rate_per_minute": 30,
          "high_entropy_data_write_peak_percent": 30,
          "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
          "newly_observed_file_extensions": [
            {
              "count": [
                20
              ],
              "name": [
                "lockile"
              ]
            }
          ],
          "time": "2021-12-01 12:46:20 -0500"
        },
        "typical_usage": {
          "file_create_peak_rate_per_minute": 50,
          "file_delete_peak_rate_per_minute": 10,
          "file_rename_peak_rate_per_minute": 5,
          "high_entropy_data_write_peak_percent": 10,
          "high_entropy_data_write_peak_rate_kb_per_minute": 1200
        }
      }
    },
    "application": {
      "name": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
    },
    "asynchronous_directory_delete": {
      "trash_bin": "string"
    },
    "autosize": {
      "mode": "string"
    },
    "clone": {
      "inherited_physical_used": 0,
      "inherited_savings": 0,
      "parent_snapshot": {
```

```
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "this_snapshot",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "parent_svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "parent_volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      },
      "split_complete_percent": 0,
      "split_estimate": 0
    },
    "cloud_retrieval_policy": "string",
    "comment": "string",
    "consistency_group": {
      "name": "consistency_group_1",
      "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": "string"
        },
        "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
          },
          "max_size": "string",
          "metadata": 0,
          "over_provisioned": 0,
          "performance_tier_footprint": 0,
          "snapshot": {
            "used": 0
          },
          "total_footprint": 0,
          "total_metadata": 0,
          "total_metadata_footprint": 0,
          "used": 0
        }
      }
    ],
    "create_time": "2018-06-04 15:00:00 -0400",
    "efficiency": {
      "application_io_size": "string",
      "auto_state": "string",
      "compaction": "string",
      "compression": "string",
      "compression_type": "string",
      "cross_volume_dedupe": "string",
      "dedupe": "string",
      "idcs_scanner": {
        "mode": "string",
```

```
        "operation_state": "string",
        "status": "string",
        "threshold_inactive_time": "P14D"
      },
      "last_op_begin": "string",
      "last_op_end": "string",
      "last_op_err": "string",
      "last_op_size": 0,
      "last_op_state": "string",
      "op_state": "string",
      "policy": {
        "name": "string"
      },
      "progress": "string",
      "scanner": {
        "state": "string"
      },
      "schedule": "string",
      "space_savings": {
        "compression": 0,
        "compression_percent": 0,
        "dedupe": 0,
        "dedupe_percent": 0,
        "dedupe_sharing": 0,
        "total": 0,
        "total_percent": 0
      },
      "state": "string",
      "storage_efficiency_mode": "string",
      "type": "string",
      "volume_path": "string"
    },
    "encryption": {
      "action": "string",
      "key_create_time": "2022-01-01 14:00:00 -0500",
      "key_id": "string",
      "state": "string",
      "status": {
        "code": "string",
        "message": "string"
      },
      "type": "string"
    },
    "files": {
      "used": 0
    },
```

```
      "flash_pool": {
        "cache_eligibility": "string",
        "cache_retention_priority": "string",
        "caching_policy": "string"
      },
      "flexcache_endpoint_type": "string",
      "flexgroup": {
        "name": "my_flexgroup",
        "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
      },
      "granular_data_mode": "string",
      "guarantee": {
        "type": "string"
      },
      "language": "string",
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "cloud": {
          "duration": "PT15S",
          "iops": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "latency": {
            "read": 200,
            "total": 1000,
            "write": 100
          },
          "status": "ok",
          "timestamp": "2017-01-25 06:20:13 -0500"
        },
        "duration": "PT15S",
        "flexcache": {
          "bandwidth_savings": 4096,
          "cache_miss_percent": 20,
          "duration": "PT1D",
          "status": "ok",
          "timestamp": "2017-01-25 06:20:13 -0500"
        },
        "iops": {
          "read": 200,
```

```
          "total": 1000,
          "write": 100
        },
        "latency": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "status": "ok",
        "throughput": {
          "read": 200,
          "total": 1000,
          "write": 100
        },
        "timestamp": "2017-01-25 06:20:13 -0500"
      },
      "movement": {
        "capacity_tier_optimized": true,
        "cutover_window": 30,
        "destination_aggregate": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "percent_complete": 0,
        "start_time": "2020-12-07 03:45:12 -0500",
        "state": "replicating",
        "tiering_policy": "string"
      },
      "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": "string",
      "unix_permissions": 493
    },
    "optimize_aggregates": true,
    "qos": {
      "policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "max_throughput_iops": 10000,
        "max_throughput_mbps": 500,
        "min_throughput_iops": 2000,
        "min_throughput_mbps": 500,
        "name": "performance",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "quota": {
      "state": "string"
    },
    "rebalancing": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "data_moved": 0,
      "engine": {
        "movement": {
          "file_moves_started": 0,
          "last_error": {
            "code": 0,
            "destination": 0,
            "file_id": 0,
            "time": "2018-06-04 15:00:00 -0400"
          },
```

```
      "most_recent_start_time": "2018-06-04 15:00:00 -0400"
    },
    "scanner": {
      "blocks_scanned": 0,
      "blocks_skipped": {
        "efficiency_blocks": 0,
        "efficiency_percent": 0,
        "fast_truncate": 0,
        "footprint_invalid": 0,
        "in_snapshot": 0,
        "incompatible": 0,
        "metadata": 0,
        "on_demand_destination": 0,
        "other": 0,
        "remote_cache": 0,
        "too_large": 0,
        "too_small": 0,
        "write_fenced": 0
      },
      "files_scanned": 0,
      "files_skipped": {
        "efficiency_blocks": 0,
        "efficiency_percent": 0,
        "fast_truncate": 0,
        "footprint_invalid": 0,
        "in_snapshot": 0,
        "incompatible": 0,
        "metadata": 0,
        "on_demand_destination": 0,
        "other": 0,
        "remote_cache": 0,
        "too_large": 0,
        "too_small": 0,
        "write_fenced": 0
      }
    }
  },
  "imbalance_percent": 0,
  "imbalance_size": 0,
  "max_constituent_imbalance_percent": 0,
  "max_runtime": "string",
  "notices": [
    {
      "arguments": [
        {
          "code": "string",
```

```
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      }
    ],
    "runtime": "string",
    "start_time": "string",
    "state": "rebalancing",
    "stop_time": "string",
    "target_used": 0,
    "used_for_imbalance": 0
  },
  "scheduled_snapshot_naming_scheme": "string",
  "snaplock": {
    "append_mode_enabled": "",
    "autocommit_period": "P30M",
    "compliance_clock_time": "2018-06-04 15:00:00 -0400",
    "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": {
    "auto_adaptive_compression_footprint_data_reduction": 0,
    "available": 0,
    "block_storage_inactive_user_data": 0,
    "block_storage_inactive_user_data_percent": 0,
```

```
        "capacity_tier_footprint": 0,
        "capacity_tier_footprint_data_reduction": 0,
        "compaction_footprint_data_reduction": 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,
        "effective_total_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
        },
        "max_size": "string",
        "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": {
            "commitment": "string",
            "defer_delete": "string",
            "delete_order": "string",
            "prefix": "string",
            "trigger": "string"
          },
          "autodelete_trigger": "string",
          "reserve_available": 0,
          "reserve_size": 0,
          "space_used_percent": 0,
          "used": 0
        },
        "snapshot_reserve_unusable": 0,
        "snapshot_spill": 0,
        "total_footprint": 0,
```

```
      "total_metadata": 0,
      "total_metadata_footprint": 0,
      "used": 0,
      "user_data": 0,
      "volume_guarantee_footprint": 0
    },
    "state": "string",
    "statistics": {
      "cifs_ops_raw": {
        "access": {
          "count": 1000,
          "total_time": 200
        },
        "audit": {
          "count": 1000,
          "total_time": 200
        },
        "create": {
          "dir": {
            "count": 1000,
            "total_time": 200
          },
          "file": {
            "count": 1000,
            "total_time": 200
          },
          "other": {
            "count": 1000,
            "total_time": 200
          },
          "symlink": {
            "count": 1000,
            "total_time": 200
          }
        },
        "getattr": {
          "count": 1000,
          "total_time": 200
        },
        "link": {
          "count": 1000,
          "total_time": 200
        },
        "lock": {
          "count": 1000,
          "total_time": 200
```

```
        },
        "lookup": {
          "count": 1000,
          "total_time": 200
        },
        "open": {
          "count": 1000,
          "total_time": 200
        },
        "read": {
          "count": 1000,
          "total_time": 200,
          "volume_protocol_latency_histogram_counts": [
            0,
            0,
            0,
            0,
            0,
            15,
            35,
            100,
            200,
            200,
            300,
            500,
            500,
            500,
            1000,
            1000,
            800,
            500,
            500,
            300,
            200,
            50,
            40,
            15,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
```

```
            0,
            0,
            0,
            0,
            0,
            0,
            0
          ],
          "volume_protocol_latency_histogram_labels": [
            "<2us",
            "<6us",
            "<10us",
            "<14us",
            "<20us",
            "<40us",
            "<60us",
            "<80us",
            "<100us",
            "<200us",
            "<400us",
            "<600us",
            "<800us",
            "<1ms",
            "<2ms",
            "<4ms",
            "<6ms",
            "<8ms",
            "<10ms",
            "<12ms",
            "<14ms",
            "<16ms",
            "<18ms",
            "<20ms",
            "<40ms",
            "<60ms",
            "<80ms",
            "<100ms",
            "<200ms",
            "<400ms",
            "<600ms",
            "<800ms",
            "<1s",
            "<2s",
            "<4s",
            "<6s",
            "<8s",
```

```
            "<10s",
            "<20s",
            ">20s"
          ],
          "volume_protocol_size_histogram_counts": [
            2400,
            1055,
            1100,
            700,
            500,
            300,
            200,
            100,
            100,
            50,
            50,
            75,
            25,
            0,
            0
          ],
          "volume_protocol_size_histogram_labels": [
            "<    4KB",
            "=    4KB",
            "<    8KB",
            "=    8KB",
            "<   16KB",
            "=   16KB",
            "<   32KB",
            "=   32KB",
            "<   64KB",
            "=   64KB",
            "<  256KB",
            "=  256KB",
            "< 1024KB",
            "= 1024KB",
            "> 1024KB"
          ]
        },
        "readdir": {
          "count": 1000,
          "total_time": 200
        },
        "readlink": {
          "count": 1000,
          "total_time": 200
```

```
        },
        "rename": {
          "count": 1000,
          "total_time": 200
        },
        "setattr": {
          "count": 1000,
          "total_time": 200
        },
        "unlink": {
          "count": 1000,
          "total_time": 200
        },
        "watch": {
          "count": 1000,
          "total_time": 200
        },
        "write": {
          "count": 1000,
          "total_time": 200,
          "volume_protocol_latency_histogram_counts": [
            0,
            0,
            0,
            0,
            0,
            15,
            35,
            100,
            200,
            200,
            300,
            500,
            500,
            500,
            1000,
            1000,
            800,
            500,
            500,
            300,
            200,
            50,
            40,
            15,
            0,
```

```
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0
          ],
          "volume_protocol_latency_histogram_labels": [
            "<2us",
            "<6us",
            "<10us",
            "<14us",
            "<20us",
            "<40us",
            "<60us",
            "<80us",
            "<100us",
            "<200us",
            "<400us",
            "<600us",
            "<800us",
            "<1ms",
            "<2ms",
            "<4ms",
            "<6ms",
            "<8ms",
            "<10ms",
            "<12ms",
            "<14ms",
            "<16ms",
            "<18ms",
            "<20ms",
            "<40ms",
            "<60ms",
            "<80ms",
            "<100ms",
            "<200ms",
```

```
            "<400ms",
            "<600ms",
            "<800ms",
            "<1s",
            "<2s",
            "<4s",
            "<6s",
            "<8s",
            "<10s",
            "<20s",
            ">20s"
          ],
          "volume_protocol_size_histogram_counts": [
            2400,
            1055,
            1100,
            700,
            500,
            300,
            200,
            100,
            100,
            50,
            50,
            75,
            25,
            0,
            0
          ],
          "volume_protocol_size_histogram_labels": [
            "<    4KB",
            "=    4KB",
            "<    8KB",
            "=    8KB",
            "<   16KB",
            "=   16KB",
            "<   32KB",
            "=   32KB",
            "<   64KB",
            "=   64KB",
            "<  256KB",
            "=  256KB",
            "< 1024KB",
            "= 1024KB",
            "> 1024KB"
          ]
```

```
      }
    },
    "cloud": {
      "iops_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "timestamp": "2017-01-25 06:20:13 -0500"
    },
    "flexcache_raw": {
      "cache_miss_blocks": 10,
      "client_requested_blocks": 500,
      "status": "ok",
      "timestamp": "2017-01-25 06:20:13 -0500"
    },
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "nfs_ops_raw": {
      "access": {
        "count": 1000,
        "total_time": 200
      },
      "audit": {
        "count": 1000,
        "total_time": 200
      },
      "create": {
        "dir": {
          "count": 1000,
          "total_time": 200
        },
```

```
            "file": {
              "count": 1000,
              "total_time": 200
            },
            "other": {
              "count": 1000,
              "total_time": 200
            },
            "symlink": {
              "count": 1000,
              "total_time": 200
            }
          },
          "getattr": {
            "count": 1000,
            "total_time": 200
          },
          "link": {
            "count": 1000,
            "total_time": 200
          },
          "lock": {
            "count": 1000,
            "total_time": 200
          },
          "lookup": {
            "count": 1000,
            "total_time": 200
          },
          "open": {
            "count": 1000,
            "total_time": 200
          },
          "read": {
            "count": 1000,
            "total_time": 200,
            "volume_protocol_latency_histogram_counts": [
              0,
              0,
              0,
              0,
              0,
              15,
              35,
              100,
              200,
```

```
            200,
            300,
            500,
            500,
            500,
            1000,
            1000,
            800,
            500,
            500,
            300,
            200,
            50,
            40,
            15,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0
          ],
          "volume_protocol_latency_histogram_labels": [
            "<2us",
            "<6us",
            "<10us",
            "<14us",
            "<20us",
            "<40us",
            "<60us",
            "<80us",
            "<100us",
            "<200us",
            "<400us",
            "<600us",
            "<800us",
```

```
            "<1ms",
            "<2ms",
            "<4ms",
            "<6ms",
            "<8ms",
            "<10ms",
            "<12ms",
            "<14ms",
            "<16ms",
            "<18ms",
            "<20ms",
            "<40ms",
            "<60ms",
            "<80ms",
            "<100ms",
            "<200ms",
            "<400ms",
            "<600ms",
            "<800ms",
            "<1s",
            "<2s",
            "<4s",
            "<6s",
            "<8s",
            "<10s",
            "<20s",
            ">20s"
          ],
          "volume_protocol_size_histogram_counts": [
            2400,
            1055,
            1100,
            700,
            500,
            300,
            200,
            100,
            100,
            50,
            50,
            75,
            25,
            0,
            0
          ],
          "volume_protocol_size_histogram_labels": [
```

```
            "<    4KB",
            "=    4KB",
            "<    8KB",
            "=    8KB",
            "<   16KB",
            "=   16KB",
            "<   32KB",
            "=   32KB",
            "<   64KB",
            "=   64KB",
            "<  256KB",
            "=  256KB",
            "< 1024KB",
            "= 1024KB",
            "> 1024KB"
          ]
        },
        "readdir": {
          "count": 1000,
          "total_time": 200
        },
        "readlink": {
          "count": 1000,
          "total_time": 200
        },
        "rename": {
          "count": 1000,
          "total_time": 200
        },
        "setattr": {
          "count": 1000,
          "total_time": 200
        },
        "unlink": {
          "count": 1000,
          "total_time": 200
        },
        "watch": {
          "count": 1000,
          "total_time": 200
        },
        "write": {
          "count": 1000,
          "total_time": 200,
          "volume_protocol_latency_histogram_counts": [
            0,
```

```
            0,
            0,
            0,
            0,
            15,
            35,
            100,
            200,
            200,
            300,
            500,
            500,
            500,
            1000,
            1000,
            800,
            500,
            500,
            300,
            200,
            50,
            40,
            15,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0,
            0
          ],
          "volume_protocol_latency_histogram_labels": [
            "<2us",
            "<6us",
            "<10us",
            "<14us",
            "<20us",
```

```
            "<40us",
            "<60us",
            "<80us",
            "<100us",
            "<200us",
            "<400us",
            "<600us",
            "<800us",
            "<1ms",
            "<2ms",
            "<4ms",
            "<6ms",
            "<8ms",
            "<10ms",
            "<12ms",
            "<14ms",
            "<16ms",
            "<18ms",
            "<20ms",
            "<40ms",
            "<60ms",
            "<80ms",
            "<100ms",
            "<200ms",
            "<400ms",
            "<600ms",
            "<800ms",
            "<1s",
            "<2s",
            "<4s",
            "<6s",
            "<8s",
            "<10s",
            "<20s",
            ">20s"
          ],
          "volume_protocol_size_histogram_counts": [
            2400,
            1055,
            1100,
            700,
            500,
            300,
            200,
            100,
            100,
```

```
            50,
            50,
            75,
            25,
            0,
            0
          ],
          "volume_protocol_size_histogram_labels": [
            "<    4KB",
            "=    4KB",
            "<    8KB",
            "=    8KB",
            "<   16KB",
            "=   16KB",
            "<   32KB",
            "=   32KB",
            "<   64KB",
            "=   64KB",
            "<  256KB",
            "=  256KB",
            "< 1024KB",
            "= 1024KB",
            "> 1024KB"
          ]
        }
      },
      "status": "ok",
      "throughput_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "timestamp": "2017-01-25 06:20:13 -0500"
    },
    "status": [
      "string"
    ],
    "style": "string",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
```

```
      },
      "tiering": {
        "object_tags": [
          "string"
        ],
        "policy": "string",
        "supported": true
      },
      "type": "string",
      "use_mirrored_aggregates": true,
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ]
}
```


Error

```
Status: Default, Error
```

Name	Type	Description
error	returned_error	

Example error


```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "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. • enum: ["off", "on"] • Introduced in: 9.10 • x-nullable: true
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	

initialization

Name	Type	Description
state	string	State of the analytics file system scan.

scan_throttle_reason

Name	Type	Description
arguments	array[string]	Arguments present in the warning message encountered.
code	string	Warning code indicating why scanner throttling is reported.
message	string	A message that provides details for scan throttling.

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
files_scanned	integer	Number of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`.
initialization	initialization	
scan_progress	integer	Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`.
scan_throttle_reason	scan_throttle_reason	

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 for all existing files and directories is currently running. If this value is "initialization_paused", this means that the initialization scan is currently paused. If this value is 'unknown', this means that there was an internal error when determining the file system analytics state for the volume. • enum: ["unknown", "initializing", "initialization_paused", "off", "on"] • Introduced in: 9.8 • x-nullable: true
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.
total_files	integer	Total number of files in the volume that the file system analytics initialization scan will process. Only returned when the state is `initializing`.
unsupported_reason	unsupported_reason	

anti_ransomware_volume_attack_detection_parameters

Name	Type	Description
based_on_file_create_op_rate	boolean	Specifies whether attack detection is based on the file create operations rate. This parameter is valid only for NAS volumes.
based_on_file_delete_op_rate	boolean	Specifies whether attack detection is based on the file delete operations rate. This parameter is valid only for NAS volumes.
based_on_file_rename_op_rate	boolean	Specifies whether attack detection is based on the file rename operations rate. This parameter is valid only for NAS volumes.
based_on_high_entropy_data_rate	boolean	Specifies whether a high entropy data rate should be considered for attack detection.
based_on_never_seen_before_file_extension	boolean	Specifies whether file extensions never seen before should be considered for attack detection.
file_create_op_rate_surge_notify_percent	integer	Specifies the percentage of surge in the file create rate up to which it is considered normal behavior.
file_delete_op_rate_surge_notify_percent	integer	Specifies the percentage of surge in the file delete rate up to which it is considered normal behavior.
file_rename_op_rate_surge_notify_percent	integer	Specifies the percent of surge in the file rename rate up to which it is considered normal behavior.

Name	Type	Description
high_entropy_data_surge_notify_percent	integer	Specifies the percentage of surge in high entropy data up to which it is considered as normal behavior. For example, if the usual high entropy data rate in the volume is 5% and if this parameter is set to 100%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 10% at any time. Similarly, if this parameter is set to 400%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 25%, and so on.
never_seen_before_file_extension_count_notify_threshold	integer	Specifies the number of files found with a never seen before file extension up to which it is considered normal behavior.
never_seen_before_file_extension_duration_in_hours	integer	Specifies the duration within which the specified number of files found with never seen before file extensions is considered normal behavior.
relaxing_popular_file_extensions	boolean	Specifies whether popular file extensions should be relaxed from being treated as a suspect for the attack. Some popular file extensions are .txt, .pdf, and so on.

_links

Name	Type	Description
suspects	href	

anti_ransomware_attack_report

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

event_log

Name	Type	Description
is_enabled_on_new_file_extension_seen	boolean	Specifies whether to send an EMS when a new file extension is discovered.
is_enabled_on_snapshot_copy_creation	boolean	Specifies whether to send an EMS when a snapshot is created.

space

Name	Type	Description
snapshot_count	integer	Total number of Anti-ransomware backup snapshots.
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 snapshots.

surge_usage

Usage values of the volume’s workload during surge. This object is no longer supported use surge_statistics instead.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Peak rate of file creates per minute in the workload of the volume during surge.
file_delete_peak_rate_per_minute	integer	Peak rate of file deletes per minute in the workload of the volume during surge.
file_rename_peak_rate_per_minute	integer	Peak rate of file renames per minute in the workload of the volume during surge.
high_entropy_data_write_peak_percent	integer	Peak percentage of high entropy data writes in the volume during surge.

Name	Type	Description
high_entropy_data_write_peak_rate_kb_per_minute	integer	Peak high entropy data write rate in the volume during surge, in KBs per minute.
time	string	Timestamp at which the first surge in the volume's workload is observed.

suspect_files

Name	Type	Description
count	integer	Total number of `suspect_files.format` files observed by the Anti-ransomware analytics engine on the volume.
entropy	string	Indicates the entropy level of this file type.
format	string	File formats observed by the Anti-ransomware analytics engine on the volume.

typical_usage

Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Typical peak rate of file creates per minute in the workload of the volume.
file_delete_peak_rate_per_minute	integer	Typical peak rate of file deletes per minute in the workload of the volume.
file_rename_peak_rate_per_minute	integer	Typical peak rate of file renames per minute in the workload of the volume.
high_entropy_data_write_peak_percent	integer	Typical peak percentage of high entropy data writes in the volume.

Name	Type	Description
high_entropy_data_write_peak_rate_kb_per_minute	integer	Typical peak high entropy data write rate in the volume, in KBs per minute.

historical_statistics

Typical usage values of volume workload.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Typical peak rate of file creates per minute in the workload of the volume.
file_delete_peak_rate_per_minute	integer	Typical peak rate of file deletes per minute in the workload of the volume.
file_rename_peak_rate_per_minute	integer	Typical peak rate of file renames per minute in the workload of the volume.
high_entropy_data_write_peak_percent	integer	Typical peak percentage of high entropy data writes in the volume.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Typical peak high entropy data write rate in the volume, in KBs per minute.

newly_observed_file_extensions

Name	Type	Description
count	integer	Count of newly observed file extensions.
name	string	Name of the newly observed file extension.

surge_statistics

Usage values of the volume’s workload during surge.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Peak rate of file creates per minute in the workload of the volume during surge.
file_delete_peak_rate_per_minute	integer	Peak rate of file deletes per minute in the workload of the volume during surge.
file_rename_peak_rate_per_minute	integer	Peak rate of file renames per minute in the workload of the volume during surge.
high_entropy_data_write_peak_percent	integer	Peak percentage of high entropy data writes in the volume during surge.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Peak high entropy data write rate in the volume during surge, in KBs per minute.
time	string	Timestamp at which the first surge in the volume's workload is observed.

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Peak rate of file creates per minute in the workload of the volume during surge.
file_delete_peak_rate_per_minute	integer	Peak rate of file deletes per minute in the workload of the volume during surge.
file_rename_peak_rate_per_minute	integer	Peak rate of file renames per minute in the workload of the volume during surge.
high_entropy_data_write_peak_percent	integer	Peak percentage of high entropy data writes in the volume during surge.

Name	Type	Description
high_entropy_data_write_peak_rate_kb_per_minute	integer	Peak high entropy data write rate in the volume during surge, in KBs per minute.
newly_observed_file_extensions	array[newly_observed_file_extensions]	New file extensions observed in the volume during surge.
time	string	Timestamp at which the first surge in the volume's workload is observed.

anti_ransomware_volume_workload

Name	Type	Description
file_extension_types_count	integer	Count of types of file extensions observed in the volume.
file_extensions_observed	array[string]	File extensions observed in the volume.
historical_statistics	historical_statistics	Typical usage values of volume workload.
newly_observed_file_extensions	array[newly_observed_file_extensions]	New file extensions observed in the volume during surge.
surge_statistics	surge_statistics	Usage values of the volume's workload during surge.
surge_usage	surge_usage	Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.
typical_usage	typical_usage	Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

anti_ransomware

Anti-ransomware related information of the volume.

Name	Type	Description
attack_detection_parameters	anti_ransomware_volume_attack_detection_parameters	
attack_probability	string	Probability of a ransomware attack. `none` No files are suspected of ransomware activity. `low` A number of files are suspected of ransomware activity. `moderate` A moderate number of files are suspected of ransomware activity. `high` 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.
event_log	event_log	
space	space	

Name	Type	Description
state	string	Anti-ransomware state. `disabled` Anti-ransomware monitoring is disabled on the volume. This is the default state in a POST operation. `disable_in_progress` Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. Valid in GET operation. `dry_run` Anti-ransomware monitoring is enabled in the evaluation mode. `enabled` Anti-ransomware monitoring is active on the volume. `paused` Anti-ransomware monitoring is paused on the volume. `enable_paused` Anti-ransomware monitoring is paused on the volume from its earlier enabled state. Valid in GET operation. `dry_run_paused` Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. Valid in GET operation. 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.
surge_as_normal	boolean	Indicates whether or not to set the surge values as historical values. This field is no longer supported. Use update_baseline_from_surge instead.
surge_usage	surge_usage	Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.
suspect_files	array[suspect_files]	
typical_usage	typical_usage	Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

Name	Type	Description
update_baseline_from_surge	boolean	Sets the observed surge value as the new baseline on a volume.
workload	anti_ransomware_volume_workload	

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.

asynchronous_directory_delete

Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.

Name	Type	Description
enabled	boolean	Specifies whether asynchronous directory delete from the client is enabled on the volume.
trash_bin	string	Name of the trash bin directory. If no "trash_bin" property is specified when enabling, the default trash bin name, "._ontaptrashbin", is used.

autosize

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

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

snapshot_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

parent_svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM. This field cannot be specified in a PATCH method.
uuid	string	The unique identifier of the SVM. This field cannot be specified in a PATCH method.

parent_volume

Name	Type	Description
_links	_links	
name	string	The name of the volume. This field cannot be specified in a PATCH method.

Name	Type	Description
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6 • x-nullable: true

clone

Name	Type	Description
has_flexclone	boolean	Specifies whether this volume is a parent of any FlexClone volume.
inherited_physical_used	integer	Inherited physical used from the clone's base snapshot.
inherited_savings	integer	Inherited savings from the clone's base snapshot.
is_flexclone	boolean	Specifies if this volume is a normal FlexVol volume or FlexClone volume. This field needs to be set when creating a FlexClone volume. Valid in POST.
parent_snapshot	snapshot_reference	
parent_svm	parent_svm	SVM, applies only to SVM-scoped objects.
parent_volume	parent_volume	
split_complete_percent	integer	Percentage of FlexClone volume blocks split from its parent volume.
split_estimate	integer	Space required by the containing-aggregate to split the FlexClone volume.

Name	Type	Description
split_initiated	boolean	This field is set when a split is executed on any FlexClone volume, that is when the FlexClone volume is split from its parent FlexVol volume. Setting this field initiates a split of a FlexClone volume from a FlexVol volume. 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 volume 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 volume constituents is 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.
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 autodelete is currently enabled on this volume.
reserve_percent	integer	The space that has been set aside as a reserve for snapshot usage, in percent.
used	integer	The total space used by snapshots 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.

Name	Type	Description
block_storage_inactive_user_data	integer	The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store.
capacity_tier_footprint	integer	Space used by capacity tier for this volume in the FabricPool aggregate, in bytes.
footprint	integer	Data used for this volume in the aggregate, in bytes.
large_size_enabled	boolean	Specifies whether the support for large volumes and large files is enabled on the volume.
local_tier_footprint	integer	Space used by the local tier for this volume in the aggregate, in bytes.
logical_space	logical_space	
max_size	string	Indicates the maximum size supported for each constituent of the FlexGroup volume.
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	

Name	Type	Description
total_footprint	integer	Data and metadata used for this volume in the aggregate, in bytes.
total_metadata	integer	Space used by the total metadata in the volume, in bytes.
total_metadata_footprint	integer	Space used by the volume metadata footprint 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	
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 volume constituents is not supported.
name	string	FlexGroup volume constituent name.
space	space	

idcs_scanner

Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner.

Name	Type	Description
enabled	boolean	Specifies the administrative state of the inactive data compression scanner. Disabling inactive data compression is not allowed on Capacity optimized Flash with QAT supported platforms.
inactive_days	integer	Data blocks older than, or equal to, 'inactive_days' are picked up by the inactive data compression scanner. Valid for PATCH only. Only applicable when 'operation_state' set to 'active'.
mode	string	Specifies the mode of inactive data compression scanner. Valid for PATCH and GET.
operation_state	string	Specifies the operational state of the inactive data compression scanner. VALID for PATCH and GET. Valid options for PATCH are "idle" and "active".
status	string	Status of last inactive data compression scan on the volume.
threshold_inactive_time	string	Time interval after which inactive data compression is automatically triggered. The value is in days and is represented in the ISO-8601 format "P<num>D", for example "P3D" represents a duration of 3 days. This field is not supported on QAT supported platforms.</num>

policy

Name	Type	Description
name	string	Specifies the name of the efficiency policy. The "inline-only" and "none" policies are not supported on Capacity optimized Flash with QAT supported platform.

scanner

Name	Type	Description
compression	boolean	Start compression if scanning old data. Valid for PATCH and GET. This option is not supported for FSX/CVO platforms.
dedupe	boolean	Start deduplication if scanning old data. Valid for PATCH and GET.
scan_old_data	boolean	Indicates whether or not to scan old data. Valid for PATCH and GET.
state	string	State of the volume efficiency scanner. Valid for PATCH and GET. Valid options for PATCH are "idle" and "active".

space_savings

Name	Type	Description
compression	integer	Total disk space that is saved by compressing blocks on the referenced file system, in bytes.
compression_percent	integer	Percentage of total disk space that is saved by compressing blocks on the referenced file system.
dedupe	integer	Total disk space that is saved by deduplication and file cloning, in bytes.
dedupe_percent	integer	Percentage of total disk space that is saved by deduplication and file cloning.
dedupe_sharing	integer	Total disk space that is shared due to deduplication and file cloning.
total	integer	Total disk space saved in the volume due to deduplication, compression and file cloning, in bytes.

Name	Type	Description
total_percent	integer	Percentage of total disk space saved in the volume due to deduplication, compression and file cloning.

efficiency

Name	Type	Description
application_io_size	string	Block size to use by compression. 8k and auto are only allowed for POST. Only auto is supported on QAT supported platforms.
auto_state	string	Automatic deduplication schedule volume state. auto ‐ Volumes with auto_state set to auto start post-process deduplication automatically. deprioritized ‐ Volumes with auto_state set to deprioritized do not start post-process deduplication automatically.
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 FlexGroup volumes, where some of the constituent volumes are compaction enabled and some are disabled.

Name	Type	Description
compression	string	The system can be enabled/disabled compression. Disabling compression is not allowed on Capacity optimized Flash with QAT supported platforms. inline ‐ Data will be compressed first and written to the volume. background ‐ Data will be written to the volume and compressed later. both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run. none ‐ None mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compression enabled and some are disabled. that On volumes with container compression enabled, background compression refers to inactive data compression scan enabled on the volume.
compression_type	string	Compression type to use by compression. Valid for PATCH and GET.

Name	Type	Description
cross_volume_dedupe	string	The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. Disabling cross volume dedupe is not allowed on Capacity optimized Flash with QAT supported platforms. inline ‐ Data will be cross volume deduped first and written to the volume. background ‐ Data will be written to the volume and cross volume deduped later. both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.
dedupe	string	The system can be enabled/disabled dedupe. Disabling dedupe is not allowed on Capacity optimized Flash with QAT supported platforms. 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 FlexGroup volumes, where some of the constituent volumes are dedupe enabled and some are disabled.
has_savings	boolean	When true, indicates that the volume contains shared(deduplication, file clones) or compressed data.

Name	Type	Description
idcs_scanner	idcs_scanner	Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner.
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.
logging_enabled	boolean	When true, indicates that space savings for any newly-written data are being logged.
op_state	string	Sis status of the volume.
policy	policy	
progress	string	Sis progress of the volume.
scanner	scanner	
schedule	string	Schedule associated with volume.
space_savings	space_savings	

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 and Cloud Volumes 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 and Cloud Volumes ONTAP ‐ enabled ‐ All supported storage efficiency features for the volume are enabled. custom ‐ Read-only field currently only supported for the FSx for ONTAP and Cloud Volumes ONTAP, user-defined storage efficiency features are enabled. For other platforms ‐ enabled ‐ At least one storage efficiency feature for the volume is enabled. • enum: ["disabled", "enabled", "mixed", "custom"] • Introduced in: 9.9 • x-nullable: true
storage_efficiency_mode	string	Storage efficiency mode used by volume. This parameter is supported only on AFF platforms. There is no difference between default and efficient modes on QAT supported platforms and auto adaptive compression is set irrespective of the modes.
type	string	Sis Type of the volume.
volume_path	string	Absolute volume path of the volume.

status

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

encryption

Name	Type	Description
action	string	This field can be used to pause an ongoing rekey or conversion operation or resume a paused rekey or conversion operation. Valid in PATCH. 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.
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_create_time	string	Encryption key creation time of the volume.

Name	Type	Description
key_id	string	The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key.
rekey	boolean	If set to 'true', re-encrypts the volume with a new key. Valid in PATCH.
state	string	Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one.
status	status	
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.

Name	Type	Description
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.

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

throughput

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

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

metric

Performance numbers, such as IOPS, latency and throughput.

Name	Type	Description
_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 volume 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 dight), first digit is assumed to be zero.

policy

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

Name	Type	Description
_links	_links	
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

Name	Type	Description
min_throughput_iops	integer	Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_mbps	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
name	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
uuid	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

qos

QoS information

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

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume 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 FlexVol volumes or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVol volumes or FlexGroup volumes when the quota state is "off".
state	string	Quota state of the volume

last_error

Error information for the last failed file move on the constituent.

Name	Type	Description
code	integer	Error code of the last file move error on the constituent.
destination	integer	DSID of the destination constituent of the last file move error on the constituent.
file_id	integer	File ID of the last file move error on the constituent.
time	string	Time of the last file move error on the constituent.

movement

Properties on this constituent related to file movement.

Name	Type	Description
file_moves_started	integer	Number of file moves started on this constituent.
last_error	last_error	Error information for the last failed file move on the constituent.

Name	Type	Description
most_recent_start_time	string	Start time of the most recent file move on the constituent.

blocks_skipped

Number of blocks skipped by the scanner on this constituent due to various reasons.

Name	Type	Description
efficiency_blocks	integer	Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.
efficiency_percent	integer	Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.
fast_truncate	integer	Number of blocks skipped by the scanner on this constituent because fast truncate is currently running on files.
footprint_invalid	integer	Number of blocks skipped by the scanner on this constituent because of files with invalid space footprints.
in_snapshot	integer	Number of blocks skipped by the scanner on this constituent because of files in snapshots.
incompatible	integer	Number of blocks skipped by the scanner on this constituent because of incompatible files.
metadata	integer	Number of blocks skipped by the scanner on this constituent because of metadata files.
on_demand_destination	integer	Number of blocks skipped by the scanner on this constituent because of on demand destination files.

Name	Type	Description
other	integer	Number of blocks skipped by the scanner on this constituent for all other reasons.
remote_cache	integer	Number of blocks skipped by the scanner on this constituent because of remote caches.
too_large	integer	Number of blocks skipped by the scanner on this constituent because of files that are larger than rebalancing.max_file_size.
too_small	integer	Number of blocks skipped by the scanner on this constituent because of files that are smaller than rebalancing.min_file_size.
write_fenced	integer	Number of blocks skipped by the scanner on this constituent because of files fenced for write operations.

files_skipped

Number of files skipped by the scanner on this constituent due to various reasons.

Name	Type	Description
efficiency_blocks	integer	Number of files skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.
efficiency_percent	integer	Number of files skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.
fast_truncate	integer	Number of files skipped by the scanner on this constituent because fast truncate is currently running on the file.
footprint_invalid	integer	Number of files skipped by the scanner on this constituent because their space footprints are invalid.

Name	Type	Description
in_snapshot	integer	Number of files skipped by the scanner on this constituent because they are trapped in snapshots.
incompatible	integer	Number of files skipped by the scanner on this constituent because they are incompatible.
metadata	integer	Number of files skipped by the scanner on this constituent because they metadata files.
on_demand_destination	integer	Number of files skipped by the scanner on this constituent because they are on demand destinations.
other	integer	Number of files skipped by the scanner on this constituent for all other reasons.
remote_cache	integer	Number of files skipped by the scanner on this constituent because they are remote caches.
too_large	integer	Number of files skipped by the scanner on this constituent because they are larger than rebalancing.max_file_size.
too_small	integer	Number of files skipped by the scanner on this constituent because they are smaller than rebalancing.min_file_size.
write_fenced	integer	Number of files skipped by the scanner on this constituent because they are fenced for write operations.

scanner

Properties related to determining which files to move and where to move them to.

Name	Type	Description
blocks_scanned	integer	Number of blocks scanned on this constituent.
blocks_skipped	blocks_skipped	Number of blocks skipped by the scanner on this constituent due to various reasons.
files_scanned	integer	Number of files scanned on this constituent.
files_skipped	files_skipped	Number of files skipped by the scanner on this constituent due to various reasons.

engine

Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions.

Name	Type	Description
movement	movement	Properties on this constituent related to file movement.
scanner	scanner	Properties related to determining which files to move and where to move them to.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

rebalancing

Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

Name	Type	Description
_links	_links	
data_moved	integer	The amount of data that has been moved in or out of a constituent. A positive value represents data moving into the constituent while a negative value is data moving out of the constituent.
engine	engine	Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions.
exclude_snapshots	boolean	Specifies whether or not to exclude files that are stuck in snapshots during rebalancing operation. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "exclude_snapshots" value. Once the operation is started, any changes to the "exclude_snapshots" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "exclude_snapshots" value.
imbalance_percent	integer	Represents the percentage the volume is out of balance.
imbalance_size	integer	Represents how much the volume is out of balance, in bytes.

Name	Type	Description
max_constituent_imbalance_percent	integer	Absolute percentage of the constituent that is most out of balance. This value will update every 30 seconds when rebalancing is not active and every 10 seconds when rebalancing is active.
max_file_moves	integer	Specifies the maximum number of concurrent file moves in a volume capacity rebalancing operation on a constituent of the FlexGroup volume. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "max_file_moves" value. Once the operation is started, any changes to the "max_file_moves" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "max_file_moves" value.
max_runtime	string	This optional field specifies the maximum time a capacity rebalancing operation runs for. Once the maximum runtime has passed, the capacity rebalancing operation stops. If it is not set, the default value is 6 hours. This value cannot be updated while a capacity rebalancing operation is running. The maximum runtime 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 "P3D" represents a duration of 3 days. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively.</num></num></num></num></num>

Name	Type	Description
max_threshold	integer	Specifies the maximum imbalance percentage for FlexGroup volume constituents. When a constituent's imbalance percentage is larger than this value, files are moved from the constituent. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "max_threshold" value. Once the operation is started, any changes to the "max_threshold" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "max_threshold" value.
min_file_size	integer	Specifies the minimum file size to consider for a volume capacity rebalancing operation. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "min_file_size" value. Once the operation is started, any changes to the "min_file_size" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "min_file_size" value. The value must be a multiple of 4KB. If it is not set, the default value is 100MB. Setting "min-file-size" to less than the default value leads to more files being moved. Moved files use granular data, which may impact read/write I/O performance.

Name	Type	Description
min_threshold	integer	Specifies the minimum imbalance percentage for FlexGroup volume constituents. When a constituent's imbalance percentage is smaller than this value, files are not moved from the constituent. When a new capacity rebalancing operation is started on a FlexGroup volume, it will use the current "min_threshold" value. Once the operation is started, any changes to the "min_threshold" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "min_threshold" value.
notices	array[error]	Capacity rebalancing notice messages.
runtime	string	Duration the capacity rebalancing operation has been running.
start_time	string	Time when the current capacity rebalancing operation started, or when a future scheduled rebalancing operation begins.

Name	Type	Description
state	string	State of the volume capacity rebalancing operation. PATCH the state to "starting" to trigger the capacity rebalance operation, and include start_time to schedule rebalancing. PATCH the state to "stopping" to stop the capacity rebalance operation, or cancel a scheduled rebalancing operation. PATCH without the state with a valid start_time to modify the start_time of an existing scheduled rebalance operation. While a FlexGroup volume is rebalancing, every constituent will have a rebalancing engine that can either be scanning the filesystem for space usage and files to move, actively moving files or temporarily doing neither. If one or more constituents has a state of "rebalancing_source" or "rebalancing_dest", then files are being moved to rebalance the FlexGroup volume. If no files are being moved, more information about what the rebalancing engine is doing for each constituent is available using the "rebalancing.engine" property. The following values apply to FlexGroup volumes. not_running ‐ capacity rebalancing is not running on the volume. starting ‐ used in a PATCH operation to start a capacity rebalancing operation. rebalancing ‐ capacity rebalancing is running on the volume. paused ‐ volume capacity rebalancing is paused on the volume. stopping ‐ used in a PATCH operation to stop a capacity rebalancing operation. unknown ‐ the system was unable to determine the rebalancing state for the volume. The following values apply to FlexGroup volume constituents.

Name	Type	Description
stop_time	string	Time when the capacity rebalancing operation stopped.
target_used	integer	Represents the ideal used size of each constituent. Calculated by dividing the total FlexGroup volume used size by the number of constituents.
used_for_imbalance	integer	Represents the used size of each constituent, as determined by the rebalancing engine. Calculated by subtracting the size used by snapshots, the size of files pending deletion and the size of filesystem metadata from the volume used size.

retention

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

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

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

snaplock

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

Name	Type	Description
autocommit_period	string	Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num>
compliance_clock_time	string	This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
expiry_time	string	Expiry time of the volume. For volumes with an infinite, indefinite or unspecified SnapLock expiry time, "9999-12-31T23:59:59" is used to denote the time.
is_audit_log	boolean	Indicates if this volume has been configured as SnapLock audit log volume for the SVM .

Name	Type	Description
litigation_count	integer	Litigation count indicates the number of active legal-holds on the volume.
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.

Name	Type	Description
is_ontap	boolean	Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to an ONTAP destination. • readOnly: 1 • Introduced in: 9.9 • x-nullable: true

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 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.

Name	Type	Description
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.
used_by_snapshots	integer	Size that is logically used across all snapshots 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.

autodelete

Name	Type	Description
commitment	string	By default, snapshot autodelete does not delete snapshots locked by Snapmirror, clones of a volume, a LUN, an NVMe namespace, or a file. Deletion of snapshots locked by these applications is specified using this option. The default value is try.
defer_delete	string	Allows the user to inform snapshot autodelete to defer the deletion of a specified snapshot until the end. The default value is user_created.

Name	Type	Description
delete_order	string	Specifies the order in which snapshot autodelete occurs. Ordering is done using the date and time the snapshot is created. The default value is oldest_first.
enabled	boolean	Specifies whether snapshot autodelete is currently enabled on this volume.
prefix	string	Specifies the prefix of the snapshot which if matched, is deleted last. Used with autodelete_defer_delete when used with a prefix value.
target_free_space	integer	Snapshots are deleted, one at a time, until the used volume space reaches the value specified. The default is 20% free space or 80% utilized.
trigger	string	Specifies when the system should trigger an autodelete of snapshots. When set to *volume*, autodelete is triggered based on volume fullness. When set to *snap_reserve*, autodelete is triggered based on snapshot reserve fullness. The default value is *volume*.

snapshot

Name	Type	Description
autodelete	autodelete	
autodelete_enabled	boolean	Specifies whether snapshot autodelete is currently enabled on this volume. This field will no longer be supported in a future release. Use autodelete.enabled instead.

Name	Type	Description
autodelete_trigger	string	Specifies when the system should trigger an autodelete of snapshots. When set to *volume*, autodelete is triggered based on volume fullness. When set to *snap_reserve*, autodelete is triggered based on snapshot reserve fullness. The default value is *volume*. This field will no longer be supported in a future release. Use autodelete.trigger instead.
reserve_available	integer	Size available for snapshots within the snapshot reserve, in bytes.
reserve_percent	integer	The space that has been set aside as a reserve for snapshot usage, in percent.
reserve_size	integer	Size in the volume that has been set aside as a reserve for snapshot usage, in bytes.
space_used_percent	integer	Percentage of snapshot reserve size that has been used.
used	integer	The total space used by snapshots in the volume, in bytes.

space

Name	Type	Description
afs_total	integer	Total size of AFS, excluding snap-reserve, in bytes.
auto_adaptive_compression_footprint_data_reduction	integer	Savings achieved due to Auto Adaptive Compression, in bytes.
available	integer	The available space, in bytes.
available_percent	integer	The space available, as a percent.

Name	Type	Description
block_storage_inactive_user_data	integer	The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store.
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.
capacity_tier_footprint_data_reduction	integer	Savings achieved in the space used by the capacity tier for this volume in the FabricPool aggregate, in bytes.
compaction_footprint_data_reduction	integer	Savings achieved due to Data Compaction, 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.

Name	Type	Description
effective_total_footprint	integer	Volume footprint after efficiency savings, in bytes. effective total footprint represents total footprint after deducting auto adaptive compression and compaction savings. effective-footprint includes aggregate metadata used by volume.
expected_available	integer	Size that should be available for the volume, irrespective of available size in the aggregate, in bytes.
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.
is_used_stale	boolean	Specifies if the virtual space used is stale.
large_size_enabled	boolean	Indicates if the support for large FlexVol volumes and large files is enabled on this volume. When configured to true, FlexVol volume size can reach up to 300TB and single file size can reach 128TB.

Name	Type	Description
local_tier_footprint	integer	Space used by the local tier for this volume in the aggregate, in bytes.
logical_space	logical_space	
max_size	string	Indicates the maximum size supported for the FlexVol volume or for each constituent of the FlexGroup volume.
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.
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 is effective total footprint which is equivalent to total footprint after deducting auto adaptive compression and compaction savings. Physical-used also includes aggregate metadata used by volume.
physical_used_percent	integer	Size that is physically used in the volume, as a percentage.

Name	Type	Description
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 snapshots 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 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.
total_metadata	integer	Space used by the total metadata in the volume, in bytes.
total_metadata_footprint	integer	Space used by the volume metadata footprint 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.
volume_guarantee_footprint	integer	Space reserved for future writes in the volume, in bytes.

access

Raw count and latency data for access operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

audit

Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

dir

Raw count and latency data for directory-create operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

file

Raw count and latency data for file-create operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

other

Raw count and latency data for create operations on objects other than files, directories and symlinks.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

symlink

Raw count and latency data for symlink-create operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

create

Raw count and latency data for create operations.

Name	Type	Description
dir	dir	Raw count and latency data for directory-create operations.
file	file	Raw count and latency data for file-create operations.
other	other	Raw count and latency data for create operations on objects other than files, directories and symlinks.
symlink	symlink	Raw count and latency data for symlink-create operations.

getattr

Raw count and latency data for getattr operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

link

Raw count and latency data for link operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

lock

Raw count and latency data for lock operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

lookup

Raw count and latency data for lookup operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

open

Raw count and latency data for open operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

read

Raw count and latency data for read operations, including histograms categorizing operations by size and latency.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • readOnly: 1 • x-ntap-advanced: true • Introduced in: 9.11 • x-nullable: true
volume_protocol_latency_histogram_counts	array[integer]	
volume_protocol_latency_histogram_labels	array[string]	Labels for the latency histogram, ranging from <2us to >20s.
volume_protocol_size_histogram_counts	array[integer]	
volume_protocol_size_histogram_labels	array[string]	Labels for the size histogram, ranging from <4KB to >1024KB.

readdir

Raw count and latency data for readdir operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

readlink

Raw count and latency data for readlink operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

rename

Raw count and latency data for rename operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

setattr

Raw count and latency data for setattr operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

unlink

Raw count and latency data for unlink operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

watch

Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

write

Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • readOnly: 1 • x-ntap-advanced: true • Introduced in: 9.11 • x-nullable: true
volume_protocol_latency_histogram_counts	array[integer]	
volume_protocol_latency_histogram_labels	array[string]	Labels for the latency histogram, ranging from <2us to >20s.
volume_protocol_size_histogram_counts	array[integer]	
volume_protocol_size_histogram_labels	array[string]	Labels for the size histogram, ranging from <4KB to >1024KB.

cifs_ops_raw

Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

Name	Type	Description
access	access	Raw count and latency data for access operations.
audit	audit	Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.
create	create	Raw count and latency data for create operations.
getattr	getattr	Raw count and latency data for getattr operations.
link	link	Raw count and latency data for link operations.
lock	lock	Raw count and latency data for lock operations.
lookup	lookup	Raw count and latency data for lookup operations.
open	open	Raw count and latency data for open operations.
read	read	Raw count and latency data for read operations, including histograms categorizing operations by size and latency.
readdir	readdir	Raw count and latency data for readdir operations.
readlink	readlink	Raw count and latency data for readlink operations.
rename	rename	Raw count and latency data for rename operations.
setattr	setattr	Raw count and latency data for setattr operations.
unlink	unlink	Raw count and latency data for unlink operations.

Name	Type	Description
watch	watch	Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.
write	write	Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

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.

Name	Type	Description
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.

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

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.

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

nfs_ops_raw

Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

Name	Type	Description
access	access	Raw count and latency data for access operations.
audit	audit	Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.
create	create	Raw count and latency data for create operations.
getattr	getattr	Raw count and latency data for getattr operations.

Name	Type	Description
link	link	Raw count and latency data for link operations.
lock	lock	Raw count and latency data for lock operations.
lookup	lookup	Raw count and latency data for lookup operations.
open	open	Raw count and latency data for open operations.
read	read	Raw count and latency data for read operations, including histograms categorizing operations by size and latency.
readdir	readdir	Raw count and latency data for readdir operations.
readlink	readlink	Raw count and latency data for readlink operations.
rename	rename	Raw count and latency data for rename operations.
setattr	setattr	Raw count and latency data for setattr operations.
unlink	unlink	Raw count and latency data for unlink operations.
watch	watch	Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.
write	write	Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

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
cifs_ops_raw	cifs_ops_raw	Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.
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.

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.
nfs_ops_raw	nfs_ops_raw	Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_ delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

Name	Type	Description
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
_links	_links	
name	string	The name of the SVM. This field cannot be specified in a PATCH method.
uuid	string	The unique identifier of the SVM. This field cannot be specified in a PATCH method.

tiering

Name	Type	Description
min_cooling_days	integer	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.
object_tags	array[string]	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".

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

volume

Name	Type	Description
_links	_links	
_tags	array[string]	Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings.
access_time_enabled	boolean	Indicates whether or not access time updates are enabled on the volume.
activity_tracking	activity_tracking	

Name	Type	Description
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
aggressive_readahead_mode	string	Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress.
analytics	analytics	
anti_ransomware	anti_ransomware	Anti-ransomware related information of the volume.
application	application	
asynchronous_directory_delete	asynchronous_directory_delete	Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.
autosize	autosize	
clone	clone	

Name	Type	Description
cloud_retrieval_policy	string	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.
cloud_write_enabled	boolean	Indicates whether or not cloud writes are enabled on the volume. NFS writes to this volume are sent to the cloud directly instead of the local performance tier. This feature is only available on volumes in FabricPools.
comment	string	A comment for the volume. Valid in POST or PATCH.
consistency_group	consistency_group	Consistency group the volume is part of.

Name	Type	Description
constituents	array[constituents]	FlexGroup volume constituents. FlexGroup volume constituents can be retrieved more efficiently by specifying "is_constituent=true" or "is_constituent=true&flexgroup.uuid=<flexgroup.uuid>" as query parameters.</flexgroup.uuid>
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	

Name	Type	Description
granular_data	boolean	State of granular data on the volume. This setting is `true` by default when creating an AWS S3 FlexGroup volume via a POST request and `false` by default for creating any other type of FlexGroup volume via a POST request. On FlexVol volumes, the setting is always `false`, as only FlexGroup volumes support this feature. Once enabled, this setting can only be disabled by restoring a snapshot. Earlier versions of ONTAP (pre 9.12) are not compatible with this feature. Therefore, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the setting was enabled. • Introduced in: 9.12 • x-nullable: true

Name	Type	Description
granular_data_mode	string	Mode of granular data on the volume. This setting defaults to `basic` when the `granular_data` parameter is set to `true`, but can be specified at the time of creation via a POST request. Earlier versions of ONTAP (pre 9.12) are not compatible with the `basic` setting. Therefore, when set to `basic`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `basic` mode was enabled. Earlier versions of ONTAP (pre 9.16) are not compatible with the `advanced` setting. Therefore, when set to `advanced`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `advanced` mode was enabled. • enum: ["disabled", "basic", "advanced"] • Introduced in: 9.16 • x-nullable: true
guarantee	guarantee	
is_object_store	boolean	Specifies whether the volume is provisioned for an object store server.
is_svm_root	boolean	Specifies whether the volume is a root volume of the SVM it belongs to.
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.

Name	Type	Description
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 volume constituents are not supported.
msid	integer	The volume's Master Set ID.
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 FlexGroup volumes, 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.
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume.

Name	Type	Description
rebalancing	rebalancing	Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.
scheduled_snapshot_naming_scheme	string	Naming Scheme for automatic snapshots: • create_time - Automatic snapshots 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 snapshots in the volume.
snapshot_directory_access_enabled	boolean	This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume.
snapshot_locking_enabled	boolean	Specifies whether or not snapshot copy locking is enabled on the volume.

Name	Type	Description
snapshot_policy	snapshot_policy	This is a reference to the snapshot policy.
space	space	
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	The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates or license. 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 volume. When the UDO License is installed, and no aggregates are specified, the system automatically provisions a FlexGroup volume on system selected aggregates. 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 ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroup volumes flexgroup_constituent ‐ FlexGroup volume constituents.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6 • x-nullable: true
validate_only	boolean	Validate the volume move or volume conversion operations and their parameters, without actually performing the operation.

returned_error

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

Create a volume on an SVM and storage aggregates

POST `/storage/volumes`

Introduced In: 9.6

Creates a volume on a specified SVM and storage aggregates.

Platform Specifics

- **Unified ONTAP**: POST must be used to create a volume.

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	The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202. • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
_tags	array[string]	Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings.

Name	Type	Description
activity_tracking	activity_tracking	
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
aggressive_readahead_mode	string	Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress.
analytics	analytics	
anti_ransomware	anti_ransomware	Anti-ransomware related information of the volume.
application	application	
asynchronous_directory_delete	asynchronous_directory_delete	Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.
autosize	autosize	
clone	clone	

Name	Type	Description
cloud_retrieval_policy	string	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.
cloud_write_enabled	boolean	Indicates whether or not cloud writes are enabled on the volume. NFS writes to this volume are sent to the cloud directly instead of the local performance tier. This feature is only available on volumes in FabricPools.
comment	string	A comment for the volume. Valid in POST or PATCH.
consistency_group	consistency_group	Consistency group the volume is part of.

Name	Type	Description
constituents	array[constituents]	FlexGroup volume constituents. FlexGroup volume constituents can be retrieved more efficiently by specifying "is_constituent=true" or "is_constituent=true&flexgroup.uuid=<flexgroup.uuid>" as query parameters.</flexgroup.uuid>
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. If a volume is being created on multiple aggregates and the "constituents_per_aggregate" field is not specified, the default value of the "constituents_per_aggregate" field is 4. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list, unless 'optimize_aggregates' is specified as 'true'. If the "aggregates.name" or "aggregates.uuid" is specified in a PATCH request to expand an existing FlexGroup volume, the default value of the "constituents_per_aggregate" field is 1. The volume expand operation is only supported on FlexGroup volumes.
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.

Name	Type	Description
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	
granular_data	boolean	State of granular data on the volume. This setting is `true` by default when creating an AWS S3 FlexGroup volume via a POST request and `false` by default for creating any other type of FlexGroup volume via a POST request. On FlexVol volumes, the setting is always `false`, as only FlexGroup volumes support this feature. Once enabled, this setting can only be disabled by restoring a snapshot. Earlier versions of ONTAP (pre 9.12) are not compatible with this feature. Therefore, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the setting was enabled. • Introduced in: 9.12 • x-nullable: true

Name	Type	Description
granular_data_mode	string	Mode of granular data on the volume. This setting defaults to `basic` when the `granular_data` parameter is set to `true`, but can be specified at the time of creation via a POST request. Earlier versions of ONTAP (pre 9.12) are not compatible with the `basic` setting. Therefore, when set to `basic`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `basic` mode was enabled. Earlier versions of ONTAP (pre 9.16) are not compatible with the `advanced` setting. Therefore, when set to `advanced`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `advanced` mode was enabled. • enum: ["disabled", "basic", "advanced"] • Introduced in: 9.16 • x-nullable: true
guarantee	guarantee	
is_object_store	boolean	Specifies whether the volume is provisioned for an object store server.
is_svm_root	boolean	Specifies whether the volume is a root volume of the SVM it belongs to.
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.

Name	Type	Description
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.
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 volume constituents are not supported.
msid	integer	The volume's Master Set ID.
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 FlexGroup volumes, 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	
optimize_aggregates	boolean	Specifies whether to create the constituents of the FlexGroup volume on the aggregates specified in the order they are specified, or whether the system should optimize the ordering of the aggregates. If this value is 'true', the system optimizes the ordering of the aggregates specified. If this value is false, the order of the aggregates is unchanged. The default value is 'false'.
qos	qos	QoS information

Name	Type	Description
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume.
rebalancing	rebalancing	Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.
scheduled_snapshot_naming_scheme	string	Naming Scheme for automatic snapshots: • create_time - Automatic snapshots 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 snapshots in the volume.
snapshot_directory_access_enabled	boolean	This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume.

Name	Type	Description
snapshot_locking_enabled	boolean	Specifies whether or not snapshot copy locking is enabled on the volume.
snapshot_policy	snapshot_policy	This is a reference to the snapshot policy.
space	space	
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.
status	array[string]	Describes the current status of a volume.

Name	Type	Description
style	string	The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates or license. 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 volume. When the UDO License is installed, and no aggregates are specified, the system automatically provisions a FlexGroup volume on system selected aggregates. 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 ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroup volumes flexgroup_constituent ‐ FlexGroup volume constituents.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup volume can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6 • x-nullable: true
validate_only	boolean	Validate the volume move or volume conversion operations and their parameters, without actually performing the operation.

Example request


```
{
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "access_time_enabled": true,
  "activity_tracking": {
    "state": "string",
    "unsupported_reason": {
      "code": "124518405",
      "message": "Volume activity tracking cannot be enabled on volumes
that contain LUNs."
    }
  },
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "aggressive_readahead_mode": "string",
  "analytics": {
    "files_scanned": 43002,
    "initialization": {
      "state": "string"
    },
    "scan_progress": 17,
    "scan_throttle_reason": {
      "arguments": [
        "string"
      ],
      "code": "6739881",
      "message": "The file system analytics scan running on volume
\"fsavol2\" in SVM \"vs2\" has slowed down on node \"my_node\". Reason:
Computing resources are being used by higher priority workloads."
    },
    "state": "string",
    "total_files": 101890,
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
    }
  },
```

```
  "anti_ransomware": {
    "attack_detection_parameters": {
      "file_create_op_rate_surge_notify_percent": 100
    },
    "attack_probability": "string",
    "attack_reports": [
      {
        "time": "2021-06-01 11:06:41 -0400"
      }
    ],
    "dry_run_start_time": "string",
    "space": {
      "snapshot_count": 0,
      "used": 0,
      "used_by_logs": 0,
      "used_by_snapshots": 0
    },
    "state": "string",
    "surge_usage": {
      "file_create_peak_rate_per_minute": 10,
      "file_delete_peak_rate_per_minute": 50,
      "file_rename_peak_rate_per_minute": 30,
      "high_entropy_data_write_peak_percent": 30,
      "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
      "time": "2021-12-01 12:46:20 -0500"
    },
    "suspect_files": [
      {
        "count": 0,
        "entropy": "string",
        "format": "string"
      }
    ],
    "typical_usage": {
      "file_create_peak_rate_per_minute": 50,
      "file_delete_peak_rate_per_minute": 10,
      "file_rename_peak_rate_per_minute": 5,
      "high_entropy_data_write_peak_percent": 10,
      "high_entropy_data_write_peak_rate_kb_per_minute": 1200
    },
    "workload": {
      "file_extension_types_count": 3,
      "file_extensions_observed": [
        "pdf",
        "jpeg",
        "txt"
```

```
    ],
    "historical_statistics": {
      "file_create_peak_rate_per_minute": 50,
      "file_delete_peak_rate_per_minute": 10,
      "file_rename_peak_rate_per_minute": 5,
      "high_entropy_data_write_peak_percent": 10,
      "high_entropy_data_write_peak_rate_kb_per_minute": 1200
    },
    "newly_observed_file_extensions": [
      {
        "count": [
          20
        ],
        "name": [
          "lockile"
        ]
      }
    ],
    "surge_statistics": {
      "file_create_peak_rate_per_minute": 10,
      "file_delete_peak_rate_per_minute": 50,
      "file_rename_peak_rate_per_minute": 30,
      "high_entropy_data_write_peak_percent": 30,
      "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
      "time": "2021-12-01 12:46:20 -0500"
    },
    "surge_usage": {
      "file_create_peak_rate_per_minute": 10,
      "file_delete_peak_rate_per_minute": 50,
      "file_rename_peak_rate_per_minute": 30,
      "high_entropy_data_write_peak_percent": 30,
      "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
      "newly_observed_file_extensions": [
        {
          "count": [
            20
          ],
          "name": [
            "lockile"
          ]
        }
      ],
      "time": "2021-12-01 12:46:20 -0500"
    },
    "typical_usage": {
      "file_create_peak_rate_per_minute": 50,
```

```
        "file_delete_peak_rate_per_minute": 10,
        "file_rename_peak_rate_per_minute": 5,
        "high_entropy_data_write_peak_percent": 10,
        "high_entropy_data_write_peak_rate_kb_per_minute": 1200
      }
    }
  },
  "application": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "asynchronous_directory_delete": {
    "trash_bin": "string"
  },
  "autosize": {
    "mode": "string"
  },
  "clone": {
    "inherited_physical_used": 0,
    "inherited_savings": 0,
    "parent_snapshot": {
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "parent_volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "split_complete_percent": 0,
    "split_estimate": 0,
    "split_initiated": true
  },
  "cloud_retrieval_policy": "string",
  "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": {
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "percent_complete": 0,
        "state": "replicating",
        "tiering_policy": "string"
      },
      "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
        },
        "max_size": "string",
        "metadata": 0,
        "over_provisioned": 0,
        "performance_tier_footprint": 0,
        "snapshot": {
          "used": 0
        },
        "total_footprint": 0,
        "total_metadata": 0,
        "total_metadata_footprint": 0,
        "used": 0
      }
    }
  ],
  "create_time": "2018-06-04 15:00:00 -0400",
  "efficiency": {
    "application_io_size": "string",
    "auto_state": "string",
    "compaction": "string",
    "compression": "string",
    "compression_type": "string",
    "cross_volume_dedupe": "string",
    "dedupe": "string",
```

```
    "idcs_scanner": {
      "mode": "string",
      "operation_state": "string",
      "status": "string",
      "threshold_inactive_time": "P14D"
    },
    "last_op_begin": "string",
    "last_op_end": "string",
    "last_op_err": "string",
    "last_op_size": 0,
    "last_op_state": "string",
    "op_state": "string",
    "policy": {
      "name": "string"
    },
    "progress": "string",
    "scanner": {
      "compression": true,
      "dedupe": true,
      "scan_old_data": true
    },
    "schedule": "string",
    "space_savings": {
      "compression": 0,
      "compression_percent": 0,
      "dedupe": 0,
      "dedupe_percent": 0,
      "dedupe_sharing": 0,
      "total": 0,
      "total_percent": 0
    },
    "state": "string",
    "storage_efficiency_mode": "string",
    "type": "string",
    "volume_path": "string"
  },
  "encryption": {
    "action": "string",
    "key_create_time": "2022-01-01 14:00:00 -0500",
    "key_id": "string",
    "key_manager_attribute": "CRN=v1:bluemix:public:containers-
kubernetes:us-south:a/asdfghjkl1234:asdfghjkl1234:worker:kubernetes-
asdfghjkl-worker1",
    "rekey": true,
    "state": "string",
    "status": {
```

```
      "code": "string",
      "message": "string"
    },
    "type": "string"
  },
  "files": {
    "used": 0
  },
  "flash_pool": {
    "cache_eligibility": "string",
    "cache_retention_priority": "string",
    "caching_policy": "string"
  },
  "flexcache_endpoint_type": "string",
  "flexgroup": {
    "name": "my_flexgroup",
    "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
  },
  "granular_data_mode": "string",
  "guarantee": {
    "type": "string"
  },
  "language": "string",
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": 0,
    "start_time": "2020-12-07 03:45:12 -0500",
    "state": "replicating",
    "tiering_policy": "string"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "id": 100,
      "name": "default"
    },
    "junction_parent": {
      "name": "vs1_root",
      "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
    },
    "path": "/user/my_volume",
    "security_style": "string",
```

```
    "unix_permissions": 493
  },
  "qos": {
    "policy": {
      "max_throughput_iops": 10000,
      "max_throughput_mbps": 500,
      "min_throughput_iops": 2000,
      "min_throughput_mbps": 500,
      "name": "performance",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "queue_for_encryption": true,
  "quota": {
    "state": "string"
  },
  "rebalancing": {
    "data_moved": 0,
    "engine": {
      "movement": {
        "file_moves_started": 0,
        "last_error": {
          "code": 0,
          "destination": 0,
          "file_id": 0,
          "time": "2018-06-04 15:00:00 -0400"
        },
        "most_recent_start_time": "2018-06-04 15:00:00 -0400"
      },
      "scanner": {
        "blocks_scanned": 0,
        "blocks_skipped": {
          "efficiency_blocks": 0,
          "efficiency_percent": 0,
          "fast_truncate": 0,
          "footprint_invalid": 0,
          "in_snapshot": 0,
          "incompatible": 0,
          "metadata": 0,
          "on_demand_destination": 0,
          "other": 0,
          "remote_cache": 0,
          "too_large": 0,
          "too_small": 0,
          "write_fenced": 0
        },
```

```
        "files_scanned": 0,
        "files_skipped": {
          "efficiency_blocks": 0,
          "efficiency_percent": 0,
          "fast_truncate": 0,
          "footprint_invalid": 0,
          "in_snapshot": 0,
          "incompatible": 0,
          "metadata": 0,
          "on_demand_destination": 0,
          "other": 0,
          "remote_cache": 0,
          "too_large": 0,
          "too_small": 0,
          "write_fenced": 0
        }
      }
    },
    "exclude_snapshots": true,
    "imbalance_percent": 0,
    "imbalance_size": 0,
    "max_constituent_imbalance_percent": 0,
    "notices": [
      {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      }
    ],
    "runtime": "string",
    "stop_time": "string",
    "target_used": 0,
    "used_for_imbalance": 0
  },
  "scheduled_snapshot_naming_scheme": "string",
  "snaplock": {
    "append_mode_enabled": "",
    "autocommit_period": "P30M",
    "compliance_clock_time": "2018-06-04 15:00:00 -0400",
    "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": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "space": {
    "auto_adaptive_compression_footprint_data_reduction": 0,
    "available": 0,
    "block_storage_inactive_user_data": 0,
    "block_storage_inactive_user_data_percent": 0,
    "capacity_tier_footprint": 0,
    "capacity_tier_footprint_data_reduction": 0,
    "compaction_footprint_data_reduction": 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,
    "effective_total_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
    },
    "max_size": "string",
    "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": {
      "reserve_available": 0,
      "reserve_size": 0,
      "space_used_percent": 0,
      "used": 0
    },
    "snapshot_reserve_unusable": 0,
    "snapshot_spill": 0,
    "total_footprint": 0,
    "total_metadata": 0,
    "total_metadata_footprint": 0,
    "used": 0,
    "user_data": 0,
    "volume_guarantee_footprint": 0
  },
  "state": "string",
  "status": [
    "string"
  ],
  "style": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "object_tags": [
      "string"
    ],
    "policy": "string"
  },
  "type": "string",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
```


Response

```
Status: 202, Accepted
```

Name	Type	Description
job	job_link	

Example response


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


Headers

Name	Description	Type
Location	Useful for tracking the resource location	string

Response

```
Status: 201, Created
```

Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
460770	[Job "jobid"] Job failed.
524561	Unable to get the information for the specified SVM's root volume because of the specified reason.
524597	Failed to create the clone volume on the node.
524597	Failed to create clone volume on the node.
524601	Parent volume "name" not found.
524819	Failed to create the clone volume.
787140	One of "aggregates.uuid", "aggregates.name", or "style" must be provided.
787141	The specified "aggregates.name" and "aggregates.uuid" refer to different aggregates.
917508	The target aggregate could not be found.
917526	The volume name specified is a duplicate.

Error Code	Description
917534	The value specified for field "-size" is too small. Update the field "-size" with the minimum size allowed and retry.
917536	Job completed with specified error.
917551	Specified snapshot policy not found.
917597	Provisioning a volume on a root aggregate is not supported.
917775	Volume \"{0}:{1}\" is not online.
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.
917886	Volume name is too long. It must be 203 characters or less.
917887	The first character of the volume name must be a letter or underscore.
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.
918232	The specified volume identifier fields must be provided.
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".

Error Code	Description
918243	The specified snapshot UUID is not correct for the specified snapshot name.
918244	Invalid "volume.type" for clone volume.
918247	Specifying a value is not valid for a volume FlexClone creation.
918252	"nas.path" is invalid.
918271	Failed to get valid export policy.
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.
918334	Cache cloud retrieval policy requires an effective cluster version of 9.14 or later.
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.
918652	"error" is an invalid value for field "-state". Valid values are "online", "offline" and "restricted".
918701	The specified operation on the volume endpoint is not supported on this platform.
1638400	Failed to retrieve snapshot information.
1638587	The specified snapshot_policy.name and snapshot_policy.uuid refer to different snapshot policies.
1638593	Operation failed because "snapdb" is disabled.
1638624	Internal error. Failed to lookup snapshot tags for \"{0}\".
2621462	The target SVM does not exist.
2621706	The specified "svm.uuid" and "svm.name" do not refer to the same SVM.
2621707	No SVM was specified. Either "svm.name" or "svm.uuid" must be supplied.
6881654	Disabling storage efficiency features is not allowed on this platform.
13107245	The specified FlexGroup volume layout exceeds the maximum number of volumes on node.

Error Code	Description
13107307	specifying this parameter when creating a flexible volume is not supported.
13107326	Request to provision FlexGroup volume "name" failed because the requested size is greater than the maximum size.
13107341	At least one valid aggregate assigned to SVM \"{0}\" with sufficient space and a homogeneous storage type is required on every node on the cluster to provision a FlexGroup volume.
13107349	Operation is only supported on flexible volumes and FlexGroup volumes.
13107405	Another volume is currently being created with the name "name" in SVM "svm.name".
13107413	Creating a FlexGroup volume is not supported on All SAN Arrays and systems that support large LUNs.
13109258	Cannot enable granular data on volume "name" in Vserver "svm.name". This setting can only be enabled on FlexGroups.
13109260	Failed to enable granular data on the volume.
13565983	A value of zero is not supported for \"-uid\" or \"-gid\".
65537463	Volume encryption keys (VEK) cannot be created or deleted for data Vserver \"{0}\". External key management has been configured for data Vserver \"{0}\" but ONTAP is not able to encrypt or decrypt with the key manager. Resolve the external key manager key issues at the key manager’s portal before creating any new encrypted volumes.
65537529	Encrypted volumes cannot be created or deleted for Vserver \"{0}\" as a rekey operation for the vserver is in progress. Try creating the encrypted volume again after some time. If the problem persists, run the rekey operation again after some time.
65537600	Encrypted volumes cannot be created or deleted for Vserver \"{0}\" while the enabled keystore configuration is being switched. If a previous attempt to switch the keystore configuration failed, or was interrupted, the system will continue to prevent encrypted volume creation for Vserver \"{0}\". Use the \"security key-manager keystore enable\" command to re-run and complete the operation.
65539430	Cannot create or delete volumes on Vserver \"{0}\" while the keystore is being initialized. Wait until the keystore is in the active state, and rerun the volume operation.

Error Code	Description
65539431	Cannot create or delete volumes on Vserver \"{0}\" while the keystore is being disabled.
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.
111411207	Volume file system analytics is not supported on volumes that contain LUNs.
111411209	Volume file system analytics is not supported on FlexCache volumes.
111411210	Volume file system analytics is not supported on audit staging volumes.
111411211	Volume file system analytics is not supported on object store server volumes.
111411212	Volume file system analytics is not supported on SnapMirror destination volumes.
111411216	Enabling or disabling volume file system analytics is not supported on individual FlexGroup constituents.
111411217	Volume file system analytics is not supported on SnapLock volumes.
111411230	Volume file system analytics is not supported on volumes that contain NVMe namespaces.
111411241	Volume file system analytics is not supported on All SAN Array clusters.
111411252	Failed to enable file system analytics on volume "name" in SVM "svm.name" because there is insufficient available space. Ensure there is at least "space.available_percent`" available space in the volume, then try the operation again.
111411253	Failed to enable file system analytics on volume "name" in SVM "svm.name" because there is insufficient available space. Ensure there is at least "space.available_percent" available space in all constituents of the FlexGroup, then try the operation again."
111411257	Failed to enable file system analytics on volume "name" in SVM "svm.name" because there is insufficient available space. Increase the volume size to "size", then try the operation again."

Error Code	Description
111415208	Cannot create a clone of volume "name" in SVM "svm.name" because the file system analytics state is not set to "off" or "on". Use REST API GET method "/api/storage/volumes/uuid?fields=analytics.state" to check the analytics state.
124518405	Volume activity tracking is not supported on volumes that contain LUNs.
124518407	Volume activity tracking is not supported on FlexCache volumes.
124518408	Volume activity tracking is not supported on audit staging volumes.
124518409	Volume activity tracking is not supported on object store server volumes.
124518410	Volume activity tracking is not supported on SnapMirror destination volumes.
124518411	Enabling or disabling volume activity tracking is not supported on individual FlexGroup constituents.
124518412	Volume activity tracking is not supported on SnapLock volumes.
124518414	Volume activity tracking is not supported on volumes that contain NVMe namespaces.
124518422	Volume activity tracking is not supported on All SAN Array clusters.

Also see the table of common errors in the Response body overview section of this documentation.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

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. • enum: ["off", "on"] • Introduced in: 9.10 • x-nullable: true
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.

Name	Type	Description
unsupported_reason	unsupported_reason	

aggregates

Aggregate

Name	Type	Description
name	string	
uuid	string	

initialization

Name	Type	Description
state	string	State of the analytics file system scan.

scan_throttle_reason

Name	Type	Description
arguments	array[string]	Arguments present in the warning message encountered.
code	string	Warning code indicating why scanner throttling is reported.
message	string	A message that provides details for scan throttling.

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
files_scanned	integer	Number of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`.
initialization	initialization	
scan_progress	integer	Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`.
scan_throttle_reason	scan_throttle_reason	
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 for all existing files and directories is currently running. If this value is "initialization_paused", this means that the initialization scan is currently paused. If this value is 'unknown', this means that there was an internal error when determining the file system analytics state for the volume. • enum: ["unknown", "initializing", "initialization_paused", "off", "on"] • Introduced in: 9.8 • x-nullable: true

Name	Type	Description
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.
total_files	integer	Total number of files in the volume that the file system analytics initialization scan will process. Only returned when the state is `initializing`.
unsupported_reason	unsupported_reason	

anti_ransomware_volume_attack_detection_parameters

Name	Type	Description
based_on_file_create_op_rate	boolean	Specifies whether attack detection is based on the file create operations rate. This parameter is valid only for NAS volumes.
based_on_file_delete_op_rate	boolean	Specifies whether attack detection is based on the file delete operations rate. This parameter is valid only for NAS volumes.
based_on_file_rename_op_rate	boolean	Specifies whether attack detection is based on the file rename operations rate. This parameter is valid only for NAS volumes.
based_on_high_entropy_data_rate	boolean	Specifies whether a high entropy data rate should be considered for attack detection.
based_on_never_seen_before_file_extension	boolean	Specifies whether file extensions never seen before should be considered for attack detection.

Name	Type	Description
file_create_op_rate_surge_notify_percent	integer	Specifies the percentage of surge in the file create rate up to which it is considered normal behavior.
file_delete_op_rate_surge_notify_percent	integer	Specifies the percentage of surge in the file delete rate up to which it is considered normal behavior.
file_rename_op_rate_surge_notify_percent	integer	Specifies the percent of surge in the file rename rate up to which it is considered normal behavior.
high_entropy_data_surge_notify_percent	integer	Specifies the percentage of surge in high entropy data up to which it is considered as normal behavior. For example, if the usual high entropy data rate in the volume is 5% and if this parameter is set to 100%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 10% at any time. Similarly, if this parameter is set to 400%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 25%, and so on.
never_seen_before_file_extension_count_notify_threshold	integer	Specifies the number of files found with a never seen before file extension up to which it is considered normal behavior.
never_seen_before_file_extension_duration_in_hours	integer	Specifies the duration within which the specified number of files found with never seen before file extensions is considered normal behavior.
relaxing_popular_file_extensions	boolean	Specifies whether popular file extensions should be relaxed from being treated as a suspect for the attack. Some popular file extensions are .txt, .pdf, and so on.

anti_ransomware_attack_report

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

event_log

Name	Type	Description
is_enabled_on_new_file_extension_seen	boolean	Specifies whether to send an EMS when a new file extension is discovered.
is_enabled_on_snapshot_copy_creation	boolean	Specifies whether to send an EMS when a snapshot is created.

space

Name	Type	Description
snapshot_count	integer	Total number of Anti-ransomware backup snapshots.
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 snapshots.

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Peak rate of file creates per minute in the workload of the volume during surge.
file_delete_peak_rate_per_minute	integer	Peak rate of file deletes per minute in the workload of the volume during surge.

Name	Type	Description
file_rename_peak_rate_per_minute	integer	Peak rate of file renames per minute in the workload of the volume during surge.
high_entropy_data_write_peak_percent	integer	Peak percentage of high entropy data writes in the volume during surge.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Peak high entropy data write rate in the volume during surge, in KBs per minute.
time	string	Timestamp at which the first surge in the volume's workload is observed.

suspect_files

Name	Type	Description
count	integer	Total number of `suspect_files.format` files observed by the Anti-ransomware analytics engine on the volume.
entropy	string	Indicates the entropy level of this file type.
format	string	File formats observed by the Anti-ransomware analytics engine on the volume.

typical_usage

Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Typical peak rate of file creates per minute in the workload of the volume.
file_delete_peak_rate_per_minute	integer	Typical peak rate of file deletes per minute in the workload of the volume.

Name	Type	Description
file_rename_peak_rate_per_minute	integer	Typical peak rate of file renames per minute in the workload of the volume.
high_entropy_data_write_peak_percent	integer	Typical peak percentage of high entropy data writes in the volume.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Typical peak high entropy data write rate in the volume, in KBs per minute.

historical_statistics

Typical usage values of volume workload.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Typical peak rate of file creates per minute in the workload of the volume.
file_delete_peak_rate_per_minute	integer	Typical peak rate of file deletes per minute in the workload of the volume.
file_rename_peak_rate_per_minute	integer	Typical peak rate of file renames per minute in the workload of the volume.
high_entropy_data_write_peak_percent	integer	Typical peak percentage of high entropy data writes in the volume.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Typical peak high entropy data write rate in the volume, in KBs per minute.

newly_observed_file_extensions

Name	Type	Description
count	integer	Count of newly observed file extensions.
name	string	Name of the newly observed file extension.

surge_statistics

Usage values of the volume's workload during surge.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Peak rate of file creates per minute in the workload of the volume during surge.
file_delete_peak_rate_per_minute	integer	Peak rate of file deletes per minute in the workload of the volume during surge.
file_rename_peak_rate_per_minute	integer	Peak rate of file renames per minute in the workload of the volume during surge.
high_entropy_data_write_peak_percent	integer	Peak percentage of high entropy data writes in the volume during surge.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Peak high entropy data write rate in the volume during surge, in KBs per minute.
time	string	Timestamp at which the first surge in the volume's workload is observed.

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Peak rate of file creates per minute in the workload of the volume during surge.
file_delete_peak_rate_per_minute	integer	Peak rate of file deletes per minute in the workload of the volume during surge.
file_rename_peak_rate_per_minute	integer	Peak rate of file renames per minute in the workload of the volume during surge.

Name	Type	Description
high_entropy_data_write_peak_percent	integer	Peak percentage of high entropy data writes in the volume during surge.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Peak high entropy data write rate in the volume during surge, in KBs per minute.
newly_observed_file_extensions	array[newly_observed_file_extensions]	New file extensions observed in the volume during surge.
time	string	Timestamp at which the first surge in the volume's workload is observed.

anti_ransomware_volume_workload

Name	Type	Description
file_extension_types_count	integer	Count of types of file extensions observed in the volume.
file_extensions_observed	array[string]	File extensions observed in the volume.
historical_statistics	historical_statistics	Typical usage values of volume workload.
newly_observed_file_extensions	array[newly_observed_file_extensions]	New file extensions observed in the volume during surge.
surge_statistics	surge_statistics	Usage values of the volume's workload during surge.
surge_usage	surge_usage	Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.
typical_usage	typical_usage	Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

anti_ransomware

Anti-ransomware related information of the volume.

Name	Type	Description
attack_detection_parameters	anti_ransomware_volume_attack_detection_parameters	
attack_probability	string	Probability of a ransomware attack. `none` No files are suspected of ransomware activity. `low` A number of files are suspected of ransomware activity. `moderate` A moderate number of files are suspected of ransomware activity. `high` 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.
event_log	event_log	
space	space	

Name	Type	Description
state	string	Anti-ransomware state. `disabled` Anti-ransomware monitoring is disabled on the volume. This is the default state in a POST operation. `disable_in_progress` Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. Valid in GET operation. `dry_run` Anti-ransomware monitoring is enabled in the evaluation mode. `enabled` Anti-ransomware monitoring is active on the volume. `paused` Anti-ransomware monitoring is paused on the volume. `enable_paused` Anti-ransomware monitoring is paused on the volume from its earlier enabled state. Valid in GET operation. `dry_run_paused` Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. Valid in GET operation. 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.
surge_as_normal	boolean	Indicates whether or not to set the surge values as historical values. This field is no longer supported. Use update_baseline_from_surge instead.
surge_usage	surge_usage	Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.
suspect_files	array[suspect_files]	
typical_usage	typical_usage	Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

Name	Type	Description
update_baseline_from_surge	boolean	Sets the observed surge value as the new baseline on a volume.
workload	anti_ransomware_volume_workload	

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.

asynchronous_directory_delete

Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.

Name	Type	Description
enabled	boolean	Specifies whether asynchronous directory delete from the client is enabled on the volume.
trash_bin	string	Name of the trash bin directory. If no "trash_bin" property is specified when enabling, the default trash bin name, "._ontaptrashbin", is used.

autosize

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

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

snapshot_reference

Name	Type	Description
name	string	
uuid	string	

parent_svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
name	string	The name of the SVM. This field cannot be specified in a PATCH method.
uuid	string	The unique identifier of the SVM. This field cannot be specified in a PATCH method.

parent_volume

Name	Type	Description
name	string	The name of the volume. This field cannot be specified in a PATCH method.

Name	Type	Description
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6 • x-nullable: true

clone

Name	Type	Description
has_flexclone	boolean	Specifies whether this volume is a parent of any FlexClone volume.
inherited_physical_used	integer	Inherited physical used from the clone's base snapshot.
inherited_savings	integer	Inherited savings from the clone's base snapshot.
is_flexclone	boolean	Specifies if this volume is a normal FlexVol volume or FlexClone volume. This field needs to be set when creating a FlexClone volume. Valid in POST.
parent_snapshot	snapshot_reference	
parent_svm	parent_svm	SVM, applies only to SVM-scoped objects.
parent_volume	parent_volume	
split_complete_percent	integer	Percentage of FlexClone volume blocks split from its parent volume.
split_estimate	integer	Space required by the containing-aggregate to split the FlexClone volume.

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 volume constituent.
uuid	string	Unique identifier for the aggregate.

destination_aggregate

Aggregate

Name	Type	Description
name	string	
uuid	string	

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup volume constituents is 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 autodelete is currently enabled on this volume.
reserve_percent	integer	The space that has been set aside as a reserve for snapshot usage, in percent.
used	integer	The total space used by snapshots 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.
large_size_enabled	boolean	Specifies whether the support for large volumes and large files is enabled on the volume.
local_tier_footprint	integer	Space used by the local tier for this volume in the aggregate, in bytes.
logical_space	logical_space	
max_size	string	Indicates the maximum size supported for each constituent of the FlexGroup volume.
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.
total_metadata	integer	Space used by the total metadata in the volume, in bytes.
total_metadata_footprint	integer	Space used by the volume metadata footprint in the aggregate, in bytes.
used	integer	The virtual space used (includes volume reserves) before storage efficiency, in bytes.

Name	Type	Description
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	
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 volume constituents is not supported.
name	string	FlexGroup volume constituent name.
space	space	

idcs_scanner

Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner.

Name	Type	Description
enabled	boolean	Specifies the administrative state of the inactive data compression scanner. Disabling inactive data compression is not allowed on Capacity optimized Flash with QAT supported platforms.
inactive_days	integer	Data blocks older than, or equal to, 'inactive_days' are picked up by the inactive data compression scanner. Valid for PATCH only. Only applicable when 'operation_state' set to 'active'.

Name	Type	Description
mode	string	Specifies the mode of inactive data compression scanner. Valid for PATCH and GET.
operation_state	string	Specifies the operational state of the inactive data compression scanner. VALID for PATCH and GET. Valid options for PATCH are "idle" and "active".
status	string	Status of last inactive data compression scan on the volume.
threshold_inactive_time	string	Time interval after which inactive data compression is automatically triggered. The value is in days and is represented in the ISO-8601 format "P<num>D", for example "P3D" represents a duration of 3 days. This field is not supported on QAT supported platforms.</num>

policy

Name	Type	Description
name	string	Specifies the name of the efficiency policy. The "inline-only" and "none" policies are not supported on Capacity optimized Flash with QAT supported platform.

scanner

space_savings

Name	Type	Description
compression	integer	Total disk space that is saved by compressing blocks on the referenced file system, in bytes.
compression_percent	integer	Percentage of total disk space that is saved by compressing blocks on the referenced file system.

Name	Type	Description
dedupe	integer	Total disk space that is saved by deduplication and file cloning, in bytes.
dedupe_percent	integer	Percentage of total disk space that is saved by deduplication and file cloning.
dedupe_sharing	integer	Total disk space that is shared due to deduplication and file cloning.
total	integer	Total disk space saved in the volume due to deduplication, compression and file cloning, in bytes.
total_percent	integer	Percentage of total disk space saved in the volume due to deduplication, compression and file cloning.

efficiency

Name	Type	Description
application_io_size	string	Block size to use by compression. 8k and auto are only allowed for POST. Only auto is supported on QAT supported platforms.
auto_state	string	Automatic deduplication schedule volume state. auto ‐ Volumes with auto_state set to auto start post-process deduplication automatically. deprioritized ‐ Volumes with auto_state set to deprioritized do not start post-process deduplication automatically.

Name	Type	Description
compaction	string	The system can be enabled/disabled compaction. inline ‐ Data will be compacted first and written to the volume. none ‐ None mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compaction enabled and some are disabled.
compression	string	The system can be enabled/disabled compression. Disabling compression is not allowed on Capacity optimized Flash with QAT supported platforms. inline ‐ Data will be compressed first and written to the volume. background ‐ Data will be written to the volume and compressed later. both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run. none ‐ None mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compression enabled and some are disabled. that On volumes with container compression enabled, background compression refers to inactive data compression scan enabled on the volume.
compression_type	string	Compression type to use by compression. Valid for PATCH and GET.

Name	Type	Description
cross_volume_dedupe	string	The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. Disabling cross volume dedupe is not allowed on Capacity optimized Flash with QAT supported platforms. inline ‐ Data will be cross volume deduped first and written to the volume. background ‐ Data will be written to the volume and cross volume deduped later. both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.
dedupe	string	The system can be enabled/disabled dedupe. Disabling dedupe is not allowed on Capacity optimized Flash with QAT supported platforms. 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 FlexGroup volumes, where some of the constituent volumes are dedupe enabled and some are disabled.
has_savings	boolean	When true, indicates that the volume contains shared(deduplication, file clones) or compressed data.

Name	Type	Description
idcs_scanner	idcs_scanner	Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner.
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.
logging_enabled	boolean	When true, indicates that space savings for any newly-written data are being logged.
op_state	string	Sis status of the volume.
policy	policy	
progress	string	Sis progress of the volume.
scanner	scanner	
schedule	string	Schedule associated with volume.
space_savings	space_savings	

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 and Cloud Volumes 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 and Cloud Volumes ONTAP ‐ enabled ‐ All supported storage efficiency features for the volume are enabled. custom ‐ Read-only field currently only supported for the FSx for ONTAP and Cloud Volumes ONTAP, user-defined storage efficiency features are enabled. For other platforms ‐ enabled ‐ At least one storage efficiency feature for the volume is enabled. • enum: ["disabled", "enabled", "mixed", "custom"] • Introduced in: 9.9 • x-nullable: true
storage_efficiency_mode	string	Storage efficiency mode used by volume. This parameter is supported only on AFF platforms. There is no difference between default and efficient modes on QAT supported platforms and auto adaptive compression is set irrespective of the modes.
type	string	Sis Type of the volume.
volume_path	string	Absolute volume path of the volume.

status

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

encryption

Name	Type	Description
action	string	This field can be used to pause an ongoing rekey or conversion operation or resume a paused rekey or conversion operation. Valid in PATCH. 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.
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_create_time	string	Encryption key creation time of the volume.

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

error_state

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

Name	Type	Description
type	string	The type of space guarantee of this volume in the aggregate.

iops

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

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

latency

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

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

cloud

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

hosted on FabricPools.

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

throughput

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

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

metric

Performance numbers, such as IOPS, latency and throughput.

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

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

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 volume constituents are not supported.

Name	Type	Description
capacity_tier_optimized	boolean	Specifies whether this move should perform optimized volume move for the capacity tier, in which data in the capacity tier does not need to be copied over.
cutover_window	integer	Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.

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

export_policy

Export Policy

Name	Type	Description
id	integer	
name	string	

junction_parent

Name	Type	Description
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.

Name	Type	Description
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 dight), first digit is assumed to be zero.

policy

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

Name	Type	Description
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

Name	Type	Description
min_throughput_iops	integer	Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_mbps	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
name	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
uuid	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

qos

QoS information

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

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume 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 FlexVol volumes or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVol volumes or FlexGroup volumes when the quota state is "off".
state	string	Quota state of the volume

last_error

Error information for the last failed file move on the constituent.

Name	Type	Description
code	integer	Error code of the last file move error on the constituent.
destination	integer	DSID of the destination constituent of the last file move error on the constituent.
file_id	integer	File ID of the last file move error on the constituent.
time	string	Time of the last file move error on the constituent.

movement

Properties on this constituent related to file movement.

Name	Type	Description
file_moves_started	integer	Number of file moves started on this constituent.
last_error	last_error	Error information for the last failed file move on the constituent.

Name	Type	Description
most_recent_start_time	string	Start time of the most recent file move on the constituent.

blocks_skipped

Number of blocks skipped by the scanner on this constituent due to various reasons.

Name	Type	Description
efficiency_blocks	integer	Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.
efficiency_percent	integer	Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.
fast_truncate	integer	Number of blocks skipped by the scanner on this constituent because fast truncate is currently running on files.
footprint_invalid	integer	Number of blocks skipped by the scanner on this constituent because of files with invalid space footprints.
in_snapshot	integer	Number of blocks skipped by the scanner on this constituent because of files in snapshots.
incompatible	integer	Number of blocks skipped by the scanner on this constituent because of incompatible files.
metadata	integer	Number of blocks skipped by the scanner on this constituent because of metadata files.
on_demand_destination	integer	Number of blocks skipped by the scanner on this constituent because of on demand destination files.

Name	Type	Description
other	integer	Number of blocks skipped by the scanner on this constituent for all other reasons.
remote_cache	integer	Number of blocks skipped by the scanner on this constituent because of remote caches.
too_large	integer	Number of blocks skipped by the scanner on this constituent because of files that are larger than rebalancing.max_file_size.
too_small	integer	Number of blocks skipped by the scanner on this constituent because of files that are smaller than rebalancing.min_file_size.
write_fenced	integer	Number of blocks skipped by the scanner on this constituent because of files fenced for write operations.

files_skipped

Number of files skipped by the scanner on this constituent due to various reasons.

Name	Type	Description
efficiency_blocks	integer	Number of files skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.
efficiency_percent	integer	Number of files skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.
fast_truncate	integer	Number of files skipped by the scanner on this constituent because fast truncate is currently running on the file.
footprint_invalid	integer	Number of files skipped by the scanner on this constituent because their space footprints are invalid.

Name	Type	Description
in_snapshot	integer	Number of files skipped by the scanner on this constituent because they are trapped in snapshots.
incompatible	integer	Number of files skipped by the scanner on this constituent because they are incompatible.
metadata	integer	Number of files skipped by the scanner on this constituent because they metadata files.
on_demand_destination	integer	Number of files skipped by the scanner on this constituent because they are on demand destinations.
other	integer	Number of files skipped by the scanner on this constituent for all other reasons.
remote_cache	integer	Number of files skipped by the scanner on this constituent because they are remote caches.
too_large	integer	Number of files skipped by the scanner on this constituent because they are larger than rebalancing.max_file_size.
too_small	integer	Number of files skipped by the scanner on this constituent because they are smaller than rebalancing.min_file_size.
write_fenced	integer	Number of files skipped by the scanner on this constituent because they are fenced for write operations.

scanner

Properties related to determining which files to move and where to move them to.

Name	Type	Description
blocks_scanned	integer	Number of blocks scanned on this constituent.
blocks_skipped	blocks_skipped	Number of blocks skipped by the scanner on this constituent due to various reasons.
files_scanned	integer	Number of files scanned on this constituent.
files_skipped	files_skipped	Number of files skipped by the scanner on this constituent due to various reasons.

engine

Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions.

Name	Type	Description
movement	movement	Properties on this constituent related to file movement.
scanner	scanner	Properties related to determining which files to move and where to move them to.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

rebalancing

Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

Name	Type	Description
data_moved	integer	The amount of data that has been moved in or out of a constituent. A positive value represents data moving into the constituent while a negative value is data moving out of the constituent.
engine	engine	Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions.
imbalance_percent	integer	Represents the percentage the volume is out of balance.
imbalance_size	integer	Represents how much the volume is out of balance, in bytes.
max_constituent_imbalance_percent	integer	Absolute percentage of the constituent that is most out of balance. This value will update every 30 seconds when rebalancing is not active and every 10 seconds when rebalancing is active.
notices	array[error]	Capacity rebalancing notice messages.
runtime	string	Duration the capacity rebalancing operation has been running.
stop_time	string	Time when the capacity rebalancing operation stopped.

Name	Type	Description
target_used	integer	Represents the ideal used size of each constituent. Calculated by dividing the total FlexGroup volume used size by the number of constituents.
used_for_imbalance	integer	Represents the used size of each constituent, as determined by the rebalancing engine. Calculated by subtracting the size used by snapshots, the size of files pending deletion and the size of filesystem metadata from the volume used size.

retention

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

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

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

snaplock

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

Name	Type	Description
autocommit_period	string	Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num>
compliance_clock_time	string	This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
expiry_time	string	Expiry time of the volume. For volumes with an infinite, indefinite or unspecified SnapLock expiry time, "9999-12-31T23:59:59" is used to denote the time.
is_audit_log	boolean	Indicates if this volume has been configured as SnapLock audit log volume for the SVM .

Name	Type	Description
litigation_count	integer	Litigation count indicates the number of active legal-holds on the volume.
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.

Name	Type	Description
is_ontap	boolean	Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to an ONTAP destination. • readOnly: 1 • Introduced in: 9.9 • x-nullable: true

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 policy.

Name	Type	Description
name	string	
uuid	string	

logical_space

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

Name	Type	Description
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.
used_by_snapshots	integer	Size that is logically used across all snapshots 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.

autodelete

Name	Type	Description
enabled	boolean	Specifies whether snapshot autodelete is currently enabled on this volume.

snapshot

Name	Type	Description
autodelete	autodelete	
autodelete_enabled	boolean	Specifies whether snapshot autodelete is currently enabled on this volume. This field will no longer be supported in a future release. Use autodelete.enabled instead.
reserve_available	integer	Size available for snapshots within the snapshot reserve, in bytes.

Name	Type	Description
reserve_percent	integer	The space that has been set aside as a reserve for snapshot usage, in percent.
reserve_size	integer	Size in the volume that has been set aside as a reserve for snapshot usage, in bytes.
space_used_percent	integer	Percentage of snapshot reserve size that has been used.
used	integer	The total space used by snapshots in the volume, in bytes.

space

Name	Type	Description
afs_total	integer	Total size of AFS, excluding snap-reserve, in bytes.
auto_adaptive_compression_footprint_data_reduction	integer	Savings achieved due to Auto Adaptive Compression, 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.

Name	Type	Description
capacity_tier_footprint_data_reduction	integer	Savings achieved in the space used by the capacity tier for this volume in the FabricPool aggregate, in bytes.
compaction_footprint_data_reduction	integer	Savings achieved due to Data Compaction, 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.
effective_total_footprint	integer	Volume footprint after efficiency savings, in bytes. effective total footprint represents total footprint after deducting auto adaptive compression and compaction savings. effective-footprint includes aggregate metadata used by volume.
expected_available	integer	Size that should be available for the volume, irrespective of available size in the aggregate, in bytes.
file_operation_metadata	integer	File operation metadata footprint, in bytes.
filesystem_size	integer	Total usable size of the volume, in bytes.

Name	Type	Description
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.
is_used_stale	boolean	Specifies if the virtual space used is stale.
large_size_enabled	boolean	Indicates if the support for large FlexVol volumes and large files is enabled on this volume. When configured to true, FlexVol volume size can reach up to 300TB and single file size can reach 128TB.
local_tier_footprint	integer	Space used by the local tier for this volume in the aggregate, in bytes.
logical_space	logical_space	
max_size	string	Indicates the maximum size supported for the FlexVol volume or for each constituent of the FlexGroup volume.
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.

Name	Type	Description
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.
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 is effective total footprint which is equivalent to total footprint after deducting auto adaptive compression and compaction savings. Physical-used also includes aggregate metadata used by volume.
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 snapshots 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 creation, in bytes.

Name	Type	Description
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.
total_metadata	integer	Space used by the total metadata in the volume, in bytes.
total_metadata_footprint	integer	Space used by the volume metadata footprint 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.
volume_guarantee_footprint	integer	Space reserved for future writes in the volume, in bytes.

access

Raw count and latency data for access operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

audit

Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

dir

Raw count and latency data for directory-create operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

file

Raw count and latency data for file-create operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

other

Raw count and latency data for create operations on objects other than files, directories and symlinks.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

symlink

Raw count and latency data for symlink-create operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

create

Raw count and latency data for create operations.

getattr

Raw count and latency data for getattr operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

link

Raw count and latency data for link operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

lock

Raw count and latency data for lock operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

lookup

Raw count and latency data for lookup operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

open

Raw count and latency data for open operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

read

Raw count and latency data for read operations, including histograms categorizing operations by size and latency.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • readOnly: 1 • x-ntap-advanced: true • Introduced in: 9.11 • x-nullable: true
volume_protocol_latency_histogram_counts	array[integer]	
volume_protocol_latency_histogram_labels	array[string]	Labels for the latency histogram, ranging from <2us to >20s.
volume_protocol_size_histogram_counts	array[integer]	
volume_protocol_size_histogram_labels	array[string]	Labels for the size histogram, ranging from <4KB to >1024KB.

readdir

Raw count and latency data for readdir operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

readlink

Raw count and latency data for readlink operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

rename

Raw count and latency data for rename operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

setattr

Raw count and latency data for setattr operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

unlink

Raw count and latency data for unlink operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

watch

Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

write

Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • readOnly: 1 • x-ntap-advanced: true • Introduced in: 9.11 • x-nullable: true
volume_protocol_latency_histogram_counts	array[integer]	
volume_protocol_latency_histogram_labels	array[string]	Labels for the latency histogram, ranging from <2us to >20s.
volume_protocol_size_histogram_counts	array[integer]	

Name	Type	Description
volume_protocol_size_histogram_labels	array[string]	Labels for the size histogram, ranging from <4KB to >1024KB.

cifs_ops_raw

Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

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.

Name	Type	Description
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.

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

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.

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

nfs_ops_raw

Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

throughput_raw

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

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

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

statistics

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

Name	Type	Description
cifs_ops_raw	cifs_ops_raw	Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.
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
nfs_ops_raw	nfs_ops_raw	Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_ delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
name	string	The name of the SVM. This field cannot be specified in a PATCH method.
uuid	string	The unique identifier of the SVM. This field cannot be specified in a PATCH method.

tiering

Name	Type	Description
min_cooling_days	integer	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.
object_tags	array[string]	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".

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

Name	Type	Description
supported	boolean	This parameter specifies whether or not FabricPools are selected when provisioning a FlexGroup volume 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 volume 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
_tags	array[string]	Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings.
activity_tracking	activity_tracking	
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.

Name	Type	Description
aggressive_readahead_mode	string	Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress.
analytics	analytics	
anti_ransomware	anti_ransomware	Anti-ransomware related information of the volume.
application	application	
asynchronous_directory_delete	asynchronous_directory_delete	Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.
autosize	autosize	
clone	clone	

Name	Type	Description
cloud_retrieval_policy	string	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.
cloud_write_enabled	boolean	Indicates whether or not cloud writes are enabled on the volume. NFS writes to this volume are sent to the cloud directly instead of the local performance tier. This feature is only available on volumes in FabricPools.
comment	string	A comment for the volume. Valid in POST or PATCH.
consistency_group	consistency_group	Consistency group the volume is part of.

Name	Type	Description
constituents	array[constituents]	FlexGroup volume constituents. FlexGroup volume constituents can be retrieved more efficiently by specifying "is_constituent=true" or "is_constituent=true&flexgroup.uuid=<flexgroup.uuid>" as query parameters.</flexgroup.uuid>
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. If a volume is being created on multiple aggregates and the "constituents_per_aggregate" field is not specified, the default value of the "constituents_per_aggregate" field is 4. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list, unless 'optimize_aggregates' is specified as 'true'. If the "aggregates.name" or "aggregates.uuid" is specified in a PATCH request to expand an existing FlexGroup volume, the default value of the "constituents_per_aggregate" field is 1. The volume expand operation is only supported on FlexGroup volumes.
convert_unicode	boolean	Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients.

Name	Type	Description
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	
granular_data	boolean	State of granular data on the volume. This setting is `true` by default when creating an AWS S3 FlexGroup volume via a POST request and `false` by default for creating any other type of FlexGroup volume via a POST request. On FlexVol volumes, the setting is always `false`, as only FlexGroup volumes support this feature. Once enabled, this setting can only be disabled by restoring a snapshot. Earlier versions of ONTAP (pre 9.12) are not compatible with this feature. Therefore, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the setting was enabled. • Introduced in: 9.12 • x-nullable: true

Name	Type	Description
granular_data_mode	string	Mode of granular data on the volume. This setting defaults to `basic` when the `granular_data` parameter is set to `true`, but can be specified at the time of creation via a POST request. Earlier versions of ONTAP (pre 9.12) are not compatible with the `basic` setting. Therefore, when set to `basic`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `basic` mode was enabled. Earlier versions of ONTAP (pre 9.16) are not compatible with the `advanced` setting. Therefore, when set to `advanced`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `advanced` mode was enabled. • enum: ["disabled", "basic", "advanced"] • Introduced in: 9.16 • x-nullable: true
guarantee	guarantee	
is_object_store	boolean	Specifies whether the volume is provisioned for an object store server.
is_svm_root	boolean	Specifies whether the volume is a root volume of the SVM it belongs to.
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.

Name	Type	Description
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.
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 volume constituents are not supported.
msid	integer	The volume's Master Set ID.
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 FlexGroup volumes, 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	
optimize_aggregates	boolean	Specifies whether to create the constituents of the FlexGroup volume on the aggregates specified in the order they are specified, or whether the system should optimize the ordering of the aggregates. If this value is 'true', the system optimizes the ordering of the aggregates specified. If this value is false, the order of the aggregates is unchanged. The default value is 'false'.

Name	Type	Description
qos	qos	QoS information
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume.
rebalancing	rebalancing	Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.
scheduled_snapshot_naming_scheme	string	Naming Scheme for automatic snapshots: • create_time - Automatic snapshots 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 snapshots in the volume.

Name	Type	Description
snapshot_directory_access_enabled	boolean	This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume.
snapshot_locking_enabled	boolean	Specifies whether or not snapshot copy locking is enabled on the volume.
snapshot_policy	snapshot_policy	This is a reference to the snapshot policy.
space	space	
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.
status	array[string]	Describes the current status of a volume.

Name	Type	Description
style	string	The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates or license. 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 volume. When the UDO License is installed, and no aggregates are specified, the system automatically provisions a FlexGroup volume on system selected aggregates. 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 ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroup volumes flexgroup_constituent ‐ FlexGroup volume constituents.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup volume can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6 • x-nullable: true
validate_only	boolean	Validate the volume move or volume conversion operations and their parameters, without actually performing the operation.

job_link

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

returned_error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

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

Delete a volume

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.

Platform Specifics

- **Unified ONTAP**: DELETE must be used to delete a volume.

Optional parameters:

- `force` - Bypasses the recovery-queue and completely removes the volume from the aggregate making it non-recoverable. By default, this flag is set to "false".

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
force	boolean	query	False	Set the force flag to "true" to bypass the recovery queue, making the deleted volume non-recoverable. • Introduced in: 9.12 • Default value:
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202. • Default value: 1 • Max value: 120 • Min value: 0

Response

```
Status: 200, Ok
```

Name	Type	Description
job	job_link	

Example response

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

Response

```
Status: 202, Accepted
```

Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
460770	[Job "jobid"] Job failed.
524615	Failed to delete volume because it has one or more clones.
917531	Cannot delete the root volume because the SVM contains other volumes.
918667	Volume "name" in SVM "svm.name" cannot be taken offline because it is configured to hold audit logs.
918693	Volume "name" in SVM "svm.name" cannot be taken offline because it is configured to hold audit logs.
918701	The specified operation on the volume endpoint is not supported on this platform.
13107349	Operation is only supported on flexible volumes and FlexGroup volumes.

Error Code	Description
65537463	Volume encryption keys (VEK) cannot be created or deleted for data Vserver \"{0}\". External key management has been configured for data Vserver \"{0}\" but ONTAP is not able to encrypt or decrypt with the key manager. Resolve the external key manager key issues at the key manager's portal before creating any new encrypted volumes.
65537529	Encrypted volumes cannot be created or deleted for Vserver \"{0}\" as a rekey operation for the vserver is in progress. Try creating the encrypted volume again after some time. If the problem persists, run the rekey operation again after some time.
65537600	Encrypted volumes cannot be created or deleted for Vserver \"{0}\" while the enabled keystore configuration is being switched. If a previous attempt to switch the keystore configuration failed, or was interrupted, the system will continue to prevent encrypted volume creation for Vserver \"{0}\". Use the \"security key-manager keystore enable\" command to re-run and complete the operation.
65539430	Cannot create or delete volumes on Vserver \"{0}\" while the keystore is being initialized. Wait until the keystore is in the active state, and rerun the volume operation.
65539431	Cannot create or delete volumes on Vserver \"{0}\" while the keystore is being disabled.

Also see the table of common errors in the Response body overview section of this documentation.

Name	Type	Description
error	returned_error	

Example error


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

returned_error

Name	Type	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 computational 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.auto_adaptive_compression_footprint_data_reduction
- space.capacity_tier_footprint_data_reduction
- space.compaction_footprint_data_reduction
- space.effective_total_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.*
- asynchronous_directory_delete.*
- rebalancing.*
- metric.*
- cloud_write_enabled

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
- volume rebalance show

- `security anti-ransomware volume show`
- `security anti-ransomware volume attack generate-report`
- `security anti-ransomware volume space show`
- `volume file async-delete client show`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Unique identifier of the volume.
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. • 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	
_tags	array[string]	Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings.
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.
aggressive_readahead_mode	string	Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress.
analytics	analytics	
anti_ransomware	anti_ransomware	Anti-ransomware related information of the volume.
application	application	
asynchronous_directory_delete	asynchronous_directory_delete	Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.
autosize	autosize	
clone	clone	

Name	Type	Description
cloud_retrieval_policy	string	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.
cloud_write_enabled	boolean	Indicates whether or not cloud writes are enabled on the volume. NFS writes to this volume are sent to the cloud directly instead of the local performance tier. This feature is only available on volumes in FabricPools.
comment	string	A comment for the volume. Valid in POST or PATCH.
consistency_group	consistency_group	Consistency group the volume is part of.

Name	Type	Description
constituents	array[constituents]	FlexGroup volume constituents. FlexGroup volume constituents can be retrieved more efficiently by specifying "is_constituent=true" or "is_constituent=true&flexgroup.uuid=<flexgroup.uuid>" as query parameters.</flexgroup.uuid>
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	

Name	Type	Description
granular_data	boolean	State of granular data on the volume. This setting is `true` by default when creating an AWS S3 FlexGroup volume via a POST request and `false` by default for creating any other type of FlexGroup volume via a POST request. On FlexVol volumes, the setting is always `false`, as only FlexGroup volumes support this feature. Once enabled, this setting can only be disabled by restoring a snapshot. Earlier versions of ONTAP (pre 9.12) are not compatible with this feature. Therefore, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the setting was enabled. • Introduced in: 9.12 • x-nullable: true

Name	Type	Description
granular_data_mode	string	Mode of granular data on the volume. This setting defaults to `basic` when the `granular_data` parameter is set to `true`, but can be specified at the time of creation via a POST request. Earlier versions of ONTAP (pre 9.12) are not compatible with the `basic` setting. Therefore, when set to `basic`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `basic` mode was enabled. Earlier versions of ONTAP (pre 9.16) are not compatible with the `advanced` setting. Therefore, when set to `advanced`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `advanced` mode was enabled. • enum: ["disabled", "basic", "advanced"] • Introduced in: 9.16 • x-nullable: true
guarantee	guarantee	
is_object_store	boolean	Specifies whether the volume is provisioned for an object store server.
is_svm_root	boolean	Specifies whether the volume is a root volume of the SVM it belongs to.
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.

Name	Type	Description
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 volume constituents are not supported.
msid	integer	The volume's Master Set ID.
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 FlexGroup volumes, 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.
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume.

Name	Type	Description
rebalancing	rebalancing	Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.
scheduled_snapshot_naming_scheme	string	Naming Scheme for automatic snapshots: • create_time - Automatic snapshots 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 snapshots in the volume.
snapshot_directory_access_enabled	boolean	This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume.
snapshot_locking_enabled	boolean	Specifies whether or not snapshot copy locking is enabled on the volume.

Name	Type	Description
snapshot_policy	snapshot_policy	This is a reference to the snapshot policy.
space	space	
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	The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates or license. 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 volume. When the UDO License is installed, and no aggregates are specified, the system automatically provisions a FlexGroup volume on system selected aggregates. 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 ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroup volumes flexgroup_constituent ‐ FlexGroup volume constituents.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6 • x-nullable: true
validate_only	boolean	Validate the volume move or volume conversion operations and their parameters, without actually performing the operation.

Example response


```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "activity_tracking": {
    "state": "string",
    "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"
    }
  ],
  "aggressive_readahead_mode": "string",
  "analytics": {
    "files_scanned": 43002,
    "initialization": {
      "state": "string"
    },
    "scan_progress": 17,
    "scan_throttle_reason": {
      "arguments": [
        "string"
      ],
      "code": "6739881",
      "message": "The file system analytics scan running on volume
\"fsavol2\" in SVM \"vs2\" has slowed down on node \"my_node\". Reason:
Computing resources are being used by higher priority workloads."
```

```
    },
    "state": "string",
    "total_files": 101890,
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
    }
  },
  "anti_ransomware": {
    "attack_detection_parameters": {
      "file_create_op_rate_surge_notify_percent": 100
    },
    "attack_probability": "string",
    "attack_reports": [
      {
        "_links": {
          "suspects": {
            "href": "/api/resourcelink"
          }
        },
        "time": "2021-06-01 11:06:41 -0400"
      }
    ],
    "dry_run_start_time": "string",
    "space": {
      "snapshot_count": 0,
      "used": 0,
      "used_by_logs": 0,
      "used_by_snapshots": 0
    },
    "state": "string",
    "surge_usage": {
      "file_create_peak_rate_per_minute": 10,
      "file_delete_peak_rate_per_minute": 50,
      "file_rename_peak_rate_per_minute": 30,
      "high_entropy_data_write_peak_percent": 30,
      "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
      "time": "2021-12-01 12:46:20 -0500"
    },
    "suspect_files": [
      {
        "count": 0,
        "entropy": "string",
        "format": "string"
      }
```

```
    ],
    "typical_usage": {
      "file_create_peak_rate_per_minute": 50,
      "file_delete_peak_rate_per_minute": 10,
      "file_rename_peak_rate_per_minute": 5,
      "high_entropy_data_write_peak_percent": 10,
      "high_entropy_data_write_peak_rate_kb_per_minute": 1200
    },
    "workload": {
      "file_extension_types_count": 3,
      "file_extensions_observed": [
        "pdf",
        "jpeg",
        "txt"
      ],
      "historical_statistics": {
        "file_create_peak_rate_per_minute": 50,
        "file_delete_peak_rate_per_minute": 10,
        "file_rename_peak_rate_per_minute": 5,
        "high_entropy_data_write_peak_percent": 10,
        "high_entropy_data_write_peak_rate_kb_per_minute": 1200
      },
      "newly_observed_file_extensions": [
        {
          "count": [
            20
          ],
          "name": [
            "lockile"
          ]
        }
      ],
      "surge_statistics": {
        "file_create_peak_rate_per_minute": 10,
        "file_delete_peak_rate_per_minute": 50,
        "file_rename_peak_rate_per_minute": 30,
        "high_entropy_data_write_peak_percent": 30,
        "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
        "time": "2021-12-01 12:46:20 -0500"
      },
      "surge_usage": {
        "file_create_peak_rate_per_minute": 10,
        "file_delete_peak_rate_per_minute": 50,
        "file_rename_peak_rate_per_minute": 30,
        "high_entropy_data_write_peak_percent": 30,
        "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
```

```
        "newly_observed_file_extensions": [
          {
            "count": [
              20
            ],
            "name": [
              "lockile"
            ]
          }
        ],
        "time": "2021-12-01 12:46:20 -0500"
      },
      "typical_usage": {
        "file_create_peak_rate_per_minute": 50,
        "file_delete_peak_rate_per_minute": 10,
        "file_rename_peak_rate_per_minute": 5,
        "high_entropy_data_write_peak_percent": 10,
        "high_entropy_data_write_peak_rate_kb_per_minute": 1200
      }
    }
  },
  "application": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "asynchronous_directory_delete": {
    "trash_bin": "string"
  },
  "autosize": {
    "mode": "string"
  },
  "clone": {
    "inherited_physical_used": 0,
    "inherited_savings": 0,
    "parent_snapshot": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "_links": {
        "self": {
```

```
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "parent_volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "split_complete_percent": 0,
    "split_estimate": 0
  },
  "cloud_retrieval_policy": "string",
  "comment": "string",
  "consistency_group": {
    "name": "consistency_group_1",
    "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": "string"
      },
      "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
        },
        "max_size": "string",
        "metadata": 0,
        "over_provisioned": 0,
        "performance_tier_footprint": 0,
        "snapshot": {
          "used": 0
        },
        "total_footprint": 0,
        "total_metadata": 0,
        "total_metadata_footprint": 0,
        "used": 0
      }
    }
  ],
  "create_time": "2018-06-04 15:00:00 -0400",
  "efficiency": {
    "application_io_size": "string",
    "auto_state": "string",
    "compaction": "string",
    "compression": "string",
    "compression_type": "string",
    "cross_volume_dedupe": "string",
    "dedupe": "string",
    "idcs_scanner": {
      "mode": "string",
      "operation_state": "string",
      "status": "string",
      "threshold_inactive_time": "P14D"
    },
    "last_op_begin": "string",
    "last_op_end": "string",
    "last_op_err": "string",
    "last_op_size": 0,
    "last_op_state": "string",
    "op_state": "string",
    "policy": {
```

```
      "name": "string"
    },
    "progress": "string",
    "scanner": {
      "state": "string"
    },
    "schedule": "string",
    "space_savings": {
      "compression": 0,
      "compression_percent": 0,
      "dedupe": 0,
      "dedupe_percent": 0,
      "dedupe_sharing": 0,
      "total": 0,
      "total_percent": 0
    },
    "state": "string",
    "storage_efficiency_mode": "string",
    "type": "string",
    "volume_path": "string"
  },
  "encryption": {
    "action": "string",
    "key_create_time": "2022-01-01 14:00:00 -0500",
    "key_id": "string",
    "state": "string",
    "status": {
      "code": "string",
      "message": "string"
    },
    "type": "string"
  },
  "files": {
    "used": 0
  },
  "flash_pool": {
    "cache_eligibility": "string",
    "cache_retention_priority": "string",
    "caching_policy": "string"
  },
  "flexcache_endpoint_type": "string",
  "flexgroup": {
    "name": "my_flexgroup",
    "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
  },
  "granular_data_mode": "string",
```

```
  "guarantee": {
    "type": "string"
  },
  "language": "string",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "cloud": {
      "duration": "PT15S",
      "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "status": "ok",
      "timestamp": "2017-01-25 06:20:13 -0500"
    },
    "duration": "PT15S",
    "flexcache": {
      "bandwidth_savings": 4096,
      "cache_miss_percent": 20,
      "duration": "PT1D",
      "status": "ok",
      "timestamp": "2017-01-25 06:20:13 -0500"
    },
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
```

```
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 06:20:13 -0500"
  },
  "movement": {
    "capacity_tier_optimized": true,
    "cutover_window": 30,
    "destination_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": 0,
    "start_time": "2020-12-07 03:45:12 -0500",
    "state": "replicating",
    "tiering_policy": "string"
  },
  "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": "string",
    "unix_permissions": 493
  },
```

```
  "optimize_aggregates": true,
  "qos": {
    "policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "max_throughput_iops": 10000,
      "max_throughput_mbps": 500,
      "min_throughput_iops": 2000,
      "min_throughput_mbps": 500,
      "name": "performance",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "quota": {
    "state": "string"
  },
  "rebalancing": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "data_moved": 0,
    "engine": {
      "movement": {
        "file_moves_started": 0,
        "last_error": {
          "code": 0,
          "destination": 0,
          "file_id": 0,
          "time": "2018-06-04 15:00:00 -0400"
        },
        "most_recent_start_time": "2018-06-04 15:00:00 -0400"
      },
      "scanner": {
        "blocks_scanned": 0,
        "blocks_skipped": {
          "efficiency_blocks": 0,
          "efficiency_percent": 0,
          "fast_truncate": 0,
          "footprint_invalid": 0,
          "in_snapshot": 0,
          "incompatible": 0,
```

```
          "metadata": 0,
          "on_demand_destination": 0,
          "other": 0,
          "remote_cache": 0,
          "too_large": 0,
          "too_small": 0,
          "write_fenced": 0
        },
        "files_scanned": 0,
        "files_skipped": {
          "efficiency_blocks": 0,
          "efficiency_percent": 0,
          "fast_truncate": 0,
          "footprint_invalid": 0,
          "in_snapshot": 0,
          "incompatible": 0,
          "metadata": 0,
          "on_demand_destination": 0,
          "other": 0,
          "remote_cache": 0,
          "too_large": 0,
          "too_small": 0,
          "write_fenced": 0
        }
      },
      "imbalance_percent": 0,
      "imbalance_size": 0,
      "max_constituent_imbalance_percent": 0,
      "max_runtime": "string",
      "notices": [
        {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
          "code": "4",
          "message": "entry doesn't exist"
        }
      ],
      "runtime": "string",
      "start_time": "string",
      "state": "rebalancing",
      "stop_time": "string",
```

```
    "target_used": 0,
    "used_for_imbalance": 0
  },
  "scheduled_snapshot_naming_scheme": "string",
  "snaplock": {
    "append_mode_enabled": "",
    "autocommit_period": "P30M",
    "compliance_clock_time": "2018-06-04 15:00:00 -0400",
    "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": {
    "auto_adaptive_compression_footprint_data_reduction": 0,
    "available": 0,
    "block_storage_inactive_user_data": 0,
    "block_storage_inactive_user_data_percent": 0,
    "capacity_tier_footprint": 0,
    "capacity_tier_footprint_data_reduction": 0,
    "compaction_footprint_data_reduction": 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,
    "effective_total_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
    },
    "max_size": "string",
    "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": {
        "commitment": "string",
        "defer_delete": "string",
        "delete_order": "string",
        "prefix": "string",
        "trigger": "string"
      },
      "autodelete_trigger": "string",
      "reserve_available": 0,
      "reserve_size": 0,
      "space_used_percent": 0,
      "used": 0
    },
    "snapshot_reserve_unusable": 0,
    "snapshot_spill": 0,
    "total_footprint": 0,
    "total_metadata": 0,
    "total_metadata_footprint": 0,
    "used": 0,
    "user_data": 0,
    "volume_guarantee_footprint": 0
  },
  "state": "string",
  "statistics": {
    "cifs_ops_raw": {
      "access": {
        "count": 1000,
```

```
        "total_time": 200
      },
      "audit": {
        "count": 1000,
        "total_time": 200
      },
      "create": {
        "dir": {
          "count": 1000,
          "total_time": 200
        },
        "file": {
          "count": 1000,
          "total_time": 200
        },
        "other": {
          "count": 1000,
          "total_time": 200
        },
        "symlink": {
          "count": 1000,
          "total_time": 200
        }
      },
      "getattr": {
        "count": 1000,
        "total_time": 200
      },
      "link": {
        "count": 1000,
        "total_time": 200
      },
      "lock": {
        "count": 1000,
        "total_time": 200
      },
      "lookup": {
        "count": 1000,
        "total_time": 200
      },
      "open": {
        "count": 1000,
        "total_time": 200
      },
      "read": {
        "count": 1000,
```

```
        "total_time": 200,
        "volume_protocol_latency_histogram_counts": [
          0,
          0,
          0,
          0,
          0,
          15,
          35,
          100,
          200,
          200,
          300,
          500,
          500,
          500,
          1000,
          1000,
          800,
          500,
          500,
          300,
          200,
          50,
          40,
          15,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0
        ],
        "volume_protocol_latency_histogram_labels": [
          "<2us",
          "<6us",
```

```
          "<10us",
          "<14us",
          "<20us",
          "<40us",
          "<60us",
          "<80us",
          "<100us",
          "<200us",
          "<400us",
          "<600us",
          "<800us",
          "<1ms",
          "<2ms",
          "<4ms",
          "<6ms",
          "<8ms",
          "<10ms",
          "<12ms",
          "<14ms",
          "<16ms",
          "<18ms",
          "<20ms",
          "<40ms",
          "<60ms",
          "<80ms",
          "<100ms",
          "<200ms",
          "<400ms",
          "<600ms",
          "<800ms",
          "<1s",
          "<2s",
          "<4s",
          "<6s",
          "<8s",
          "<10s",
          "<20s",
          ">20s"
        ],
        "volume_protocol_size_histogram_counts": [
          2400,
          1055,
          1100,
          700,
          500,
          300,
```

```
          200,
          100,
          100,
          50,
          50,
          75,
          25,
          0,
          0
        ],
        "volume_protocol_size_histogram_labels": [
          "<    4KB",
          "=    4KB",
          "<    8KB",
          "=    8KB",
          "<   16KB",
          "=   16KB",
          "<   32KB",
          "=   32KB",
          "<   64KB",
          "=   64KB",
          "<  256KB",
          "=  256KB",
          "< 1024KB",
          "= 1024KB",
          "> 1024KB"
        ]
      },
      "readdir": {
        "count": 1000,
        "total_time": 200
      },
      "readlink": {
        "count": 1000,
        "total_time": 200
      },
      "rename": {
        "count": 1000,
        "total_time": 200
      },
      "setattr": {
        "count": 1000,
        "total_time": 200
      },
      "unlink": {
        "count": 1000,
```

```
      "total_time": 200
    },
    "watch": {
      "count": 1000,
      "total_time": 200
    },
    "write": {
      "count": 1000,
      "total_time": 200,
      "volume_protocol_latency_histogram_counts": [
        0,
        0,
        0,
        0,
        0,
        15,
        35,
        100,
        200,
        200,
        300,
        500,
        500,
        500,
        1000,
        1000,
        800,
        500,
        500,
        300,
        200,
        50,
        40,
        15,
        0,
        0,
        0,
        0,
        0,
        0,
        0,
        0,
        0,
        0,
        0,
        0,
        0,
```

```
          0,
          0,
          0,
          0
        ],
        "volume_protocol_latency_histogram_labels": [
          "<2us",
          "<6us",
          "<10us",
          "<14us",
          "<20us",
          "<40us",
          "<60us",
          "<80us",
          "<100us",
          "<200us",
          "<400us",
          "<600us",
          "<800us",
          "<1ms",
          "<2ms",
          "<4ms",
          "<6ms",
          "<8ms",
          "<10ms",
          "<12ms",
          "<14ms",
          "<16ms",
          "<18ms",
          "<20ms",
          "<40ms",
          "<60ms",
          "<80ms",
          "<100ms",
          "<200ms",
          "<400ms",
          "<600ms",
          "<800ms",
          "<1s",
          "<2s",
          "<4s",
          "<6s",
          "<8s",
          "<10s",
          "<20s",
          ">20s"
```

```
        ],
        "volume_protocol_size_histogram_counts": [
          2400,
          1055,
          1100,
          700,
          500,
          300,
          200,
          100,
          100,
          50,
          50,
          75,
          25,
          0,
          0
        ],
        "volume_protocol_size_histogram_labels": [
          "<    4KB",
          "=    4KB",
          "<    8KB",
          "=    8KB",
          "<   16KB",
          "=   16KB",
          "<   32KB",
          "=   32KB",
          "<   64KB",
          "=   64KB",
          "<  256KB",
          "=  256KB",
          "< 1024KB",
          "= 1024KB",
          "> 1024KB"
        ]
      }
    },
    "cloud": {
      "iops_raw": {
        "read": 200,
        "total": 1000,
        "write": 100
      },
      "latency_raw": {
        "read": 200,
        "total": 1000,
```

```
        "write": 100
      },
      "status": "ok",
      "timestamp": "2017-01-25 06:20:13 -0500"
    },
    "flexcache_raw": {
      "cache_miss_blocks": 10,
      "client_requested_blocks": 500,
      "status": "ok",
      "timestamp": "2017-01-25 06:20:13 -0500"
    },
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "nfs_ops_raw": {
      "access": {
        "count": 1000,
        "total_time": 200
      },
      "audit": {
        "count": 1000,
        "total_time": 200
      },
      "create": {
        "dir": {
          "count": 1000,
          "total_time": 200
        },
        "file": {
          "count": 1000,
          "total_time": 200
        },
        "other": {
          "count": 1000,
          "total_time": 200
        },
        "symlink": {
          "count": 1000,
          "total_time": 200
```

```
        }
      },
      "getattr": {
        "count": 1000,
        "total_time": 200
      },
      "link": {
        "count": 1000,
        "total_time": 200
      },
      "lock": {
        "count": 1000,
        "total_time": 200
      },
      "lookup": {
        "count": 1000,
        "total_time": 200
      },
      "open": {
        "count": 1000,
        "total_time": 200
      },
      "read": {
        "count": 1000,
        "total_time": 200,
        "volume_protocol_latency_histogram_counts": [
          0,
          0,
          0,
          0,
          0,
          15,
          35,
          100,
          200,
          200,
          300,
          500,
          500,
          500,
          1000,
          1000,
          800,
          500,
          500,
          300,
```

```
          200,
          50,
          40,
          15,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0
        ],
        "volume_protocol_latency_histogram_labels": [
          "<2us",
          "<6us",
          "<10us",
          "<14us",
          "<20us",
          "<40us",
          "<60us",
          "<80us",
          "<100us",
          "<200us",
          "<400us",
          "<600us",
          "<800us",
          "<1ms",
          "<2ms",
          "<4ms",
          "<6ms",
          "<8ms",
          "<10ms",
          "<12ms",
          "<14ms",
          "<16ms",
          "<18ms",
          "<20ms",
```

```
        "<40ms",
        "<60ms",
        "<80ms",
        "<100ms",
        "<200ms",
        "<400ms",
        "<600ms",
        "<800ms",
        "<1s",
        "<2s",
        "<4s",
        "<6s",
        "<8s",
        "<10s",
        "<20s",
        ">20s"
      ],
      "volume_protocol_size_histogram_counts": [
        2400,
        1055,
        1100,
        700,
        500,
        300,
        200,
        100,
        100,
        50,
        50,
        75,
        25,
        0,
        0
      ],
      "volume_protocol_size_histogram_labels": [
        "<    4KB",
        "=    4KB",
        "<    8KB",
        "=    8KB",
        "<   16KB",
        "=   16KB",
        "<   32KB",
        "=   32KB",
        "<   64KB",
        "=   64KB",
        "<  256KB",
```

```
          "=  256KB",
          "< 1024KB",
          "= 1024KB",
          "> 1024KB"
        ]
      },
      "readdir": {
        "count": 1000,
        "total_time": 200
      },
      "readlink": {
        "count": 1000,
        "total_time": 200
      },
      "rename": {
        "count": 1000,
        "total_time": 200
      },
      "setattr": {
        "count": 1000,
        "total_time": 200
      },
      "unlink": {
        "count": 1000,
        "total_time": 200
      },
      "watch": {
        "count": 1000,
        "total_time": 200
      },
      "write": {
        "count": 1000,
        "total_time": 200,
        "volume_protocol_latency_histogram_counts": [
          0,
          0,
          0,
          0,
          0,
          15,
          35,
          100,
          200,
          200,
          300,
          500,
```

```
          500,
          500,
          1000,
          1000,
          800,
          500,
          500,
          300,
          200,
          50,
          40,
          15,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0,
          0
        ],
        "volume_protocol_latency_histogram_labels": [
          "<2us",
          "<6us",
          "<10us",
          "<14us",
          "<20us",
          "<40us",
          "<60us",
          "<80us",
          "<100us",
          "<200us",
          "<400us",
          "<600us",
          "<800us",
          "<1ms",
          "<2ms",
          "<4ms",
```

```
        "<6ms",
        "<8ms",
        "<10ms",
        "<12ms",
        "<14ms",
        "<16ms",
        "<18ms",
        "<20ms",
        "<40ms",
        "<60ms",
        "<80ms",
        "<100ms",
        "<200ms",
        "<400ms",
        "<600ms",
        "<800ms",
        "<1s",
        "<2s",
        "<4s",
        "<6s",
        "<8s",
        "<10s",
        "<20s",
        ">20s"
      ],
      "volume_protocol_size_histogram_counts": [
        2400,
        1055,
        1100,
        700,
        500,
        300,
        200,
        100,
        100,
        50,
        50,
        75,
        25,
        0,
        0
      ],
      "volume_protocol_size_histogram_labels": [
        "<    4KB",
        "=    4KB",
        "<    8KB",
```

```
          "=    8KB",
          "<   16KB",
          "=   16KB",
          "<   32KB",
          "=   32KB",
          "<   64KB",
          "=   64KB",
          "<  256KB",
          "=  256KB",
          "< 1024KB",
          "= 1024KB",
          "> 1024KB"
        ]
      }
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 06:20:13 -0500"
  },
  "status": [
    "string"
  ],
  "style": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "object_tags": [
      "string"
    ],
    "policy": "string",
    "supported": true
  },
  "type": "string",
  "use_mirrored_aggregates": true,
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
```

```
}
```

Error

```
Status: Default, Error
```

Name	Type	Description
error	returned_error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "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. • enum: ["off", "on"] • Introduced in: 9.10 • x-nullable: true

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	

initialization

Name	Type	Description
state	string	State of the analytics file system scan.

scan_throttle_reason

Name	Type	Description
arguments	array[string]	Arguments present in the warning message encountered.
code	string	Warning code indicating why scanner throttling is reported.
message	string	A message that provides details for scan throttling.

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
files_scanned	integer	Number of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`.
initialization	initialization	
scan_progress	integer	Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`.
scan_throttle_reason	scan_throttle_reason	

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 for all existing files and directories is currently running. If this value is "initialization_paused", this means that the initialization scan is currently paused. If this value is 'unknown', this means that there was an internal error when determining the file system analytics state for the volume. • enum: ["unknown", "initializing", "initialization_paused", "off", "on"] • Introduced in: 9.8 • x-nullable: true
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.
total_files	integer	Total number of files in the volume that the file system analytics initialization scan will process. Only returned when the state is `initializing`.
unsupported_reason	unsupported_reason	

anti_ransomware_volume_attack_detection_parameters

Name	Type	Description
based_on_file_create_op_rate	boolean	Specifies whether attack detection is based on the file create operations rate. This parameter is valid only for NAS volumes.
based_on_file_delete_op_rate	boolean	Specifies whether attack detection is based on the file delete operations rate. This parameter is valid only for NAS volumes.
based_on_file_rename_op_rate	boolean	Specifies whether attack detection is based on the file rename operations rate. This parameter is valid only for NAS volumes.
based_on_high_entropy_data_rate	boolean	Specifies whether a high entropy data rate should be considered for attack detection.
based_on_never_seen_before_file_extension	boolean	Specifies whether file extensions never seen before should be considered for attack detection.
file_create_op_rate_surge_notify_percent	integer	Specifies the percentage of surge in the file create rate up to which it is considered normal behavior.
file_delete_op_rate_surge_notify_percent	integer	Specifies the percentage of surge in the file delete rate up to which it is considered normal behavior.
file_rename_op_rate_surge_notify_percent	integer	Specifies the percent of surge in the file rename rate up to which it is considered normal behavior.

Name	Type	Description
high_entropy_data_surge_notify_percent	integer	Specifies the percentage of surge in high entropy data up to which it is considered as normal behavior. For example, if the usual high entropy data rate in the volume is 5% and if this parameter is set to 100%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 10% at any time. Similarly, if this parameter is set to 400%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 25%, and so on.
never_seen_before_file_extension_count_notify_threshold	integer	Specifies the number of files found with a never seen before file extension up to which it is considered normal behavior.
never_seen_before_file_extension_duration_in_hours	integer	Specifies the duration within which the specified number of files found with never seen before file extensions is considered normal behavior.
relaxing_popular_file_extensions	boolean	Specifies whether popular file extensions should be relaxed from being treated as a suspect for the attack. Some popular file extensions are .txt, .pdf, and so on.

_links

Name	Type	Description
suspects	href	

anti_ransomware_attack_report

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

event_log

Name	Type	Description
is_enabled_on_new_file_extension_seen	boolean	Specifies whether to send an EMS when a new file extension is discovered.
is_enabled_on_snapshot_copy_creation	boolean	Specifies whether to send an EMS when a snapshot is created.

space

Name	Type	Description
snapshot_count	integer	Total number of Anti-ransomware backup snapshots.
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 snapshots.

surge_usage

Usage values of the volume’s workload during surge. This object is no longer supported use surge_statistics instead.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Peak rate of file creates per minute in the workload of the volume during surge.
file_delete_peak_rate_per_minute	integer	Peak rate of file deletes per minute in the workload of the volume during surge.
file_rename_peak_rate_per_minute	integer	Peak rate of file renames per minute in the workload of the volume during surge.
high_entropy_data_write_peak_percent	integer	Peak percentage of high entropy data writes in the volume during surge.

Name	Type	Description
high_entropy_data_write_peak_rate_kb_per_minute	integer	Peak high entropy data write rate in the volume during surge, in KBs per minute.
time	string	Timestamp at which the first surge in the volume's workload is observed.

suspect_files

Name	Type	Description
count	integer	Total number of `suspect_files.format` files observed by the Anti-ransomware analytics engine on the volume.
entropy	string	Indicates the entropy level of this file type.
format	string	File formats observed by the Anti-ransomware analytics engine on the volume.

typical_usage

Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Typical peak rate of file creates per minute in the workload of the volume.
file_delete_peak_rate_per_minute	integer	Typical peak rate of file deletes per minute in the workload of the volume.
file_rename_peak_rate_per_minute	integer	Typical peak rate of file renames per minute in the workload of the volume.
high_entropy_data_write_peak_percent	integer	Typical peak percentage of high entropy data writes in the volume.

Name	Type	Description
high_entropy_data_write_peak_rate_kb_per_minute	integer	Typical peak high entropy data write rate in the volume, in KBs per minute.

historical_statistics

Typical usage values of volume workload.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Typical peak rate of file creates per minute in the workload of the volume.
file_delete_peak_rate_per_minute	integer	Typical peak rate of file deletes per minute in the workload of the volume.
file_rename_peak_rate_per_minute	integer	Typical peak rate of file renames per minute in the workload of the volume.
high_entropy_data_write_peak_percent	integer	Typical peak percentage of high entropy data writes in the volume.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Typical peak high entropy data write rate in the volume, in KBs per minute.

newly_observed_file_extensions

Name	Type	Description
count	integer	Count of newly observed file extensions.
name	string	Name of the newly observed file extension.

surge_statistics

Usage values of the volume’s workload during surge.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Peak rate of file creates per minute in the workload of the volume during surge.
file_delete_peak_rate_per_minute	integer	Peak rate of file deletes per minute in the workload of the volume during surge.
file_rename_peak_rate_per_minute	integer	Peak rate of file renames per minute in the workload of the volume during surge.
high_entropy_data_write_peak_percent	integer	Peak percentage of high entropy data writes in the volume during surge.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Peak high entropy data write rate in the volume during surge, in KBs per minute.
time	string	Timestamp at which the first surge in the volume's workload is observed.

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Peak rate of file creates per minute in the workload of the volume during surge.
file_delete_peak_rate_per_minute	integer	Peak rate of file deletes per minute in the workload of the volume during surge.
file_rename_peak_rate_per_minute	integer	Peak rate of file renames per minute in the workload of the volume during surge.
high_entropy_data_write_peak_percent	integer	Peak percentage of high entropy data writes in the volume during surge.

Name	Type	Description
high_entropy_data_write_peak_rate_kb_per_minute	integer	Peak high entropy data write rate in the volume during surge, in KBs per minute.
newly_observed_file_extensions	array[newly_observed_file_extensions]	New file extensions observed in the volume during surge.
time	string	Timestamp at which the first surge in the volume's workload is observed.

anti_ransomware_volume_workload

Name	Type	Description
file_extension_types_count	integer	Count of types of file extensions observed in the volume.
file_extensions_observed	array[string]	File extensions observed in the volume.
historical_statistics	historical_statistics	Typical usage values of volume workload.
newly_observed_file_extensions	array[newly_observed_file_extensions]	New file extensions observed in the volume during surge.
surge_statistics	surge_statistics	Usage values of the volume's workload during surge.
surge_usage	surge_usage	Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.
typical_usage	typical_usage	Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

anti_ransomware

Anti-ransomware related information of the volume.

Name	Type	Description
attack_detection_parameters	anti_ransomware_volume_attack_detection_parameters	
attack_probability	string	Probability of a ransomware attack. `none` No files are suspected of ransomware activity. `low` A number of files are suspected of ransomware activity. `moderate` A moderate number of files are suspected of ransomware activity. `high` 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.
event_log	event_log	
space	space	

Name	Type	Description
state	string	Anti-ransomware state. `disabled` Anti-ransomware monitoring is disabled on the volume. This is the default state in a POST operation. `disable_in_progress` Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. Valid in GET operation. `dry_run` Anti-ransomware monitoring is enabled in the evaluation mode. `enabled` Anti-ransomware monitoring is active on the volume. `paused` Anti-ransomware monitoring is paused on the volume. `enable_paused` Anti-ransomware monitoring is paused on the volume from its earlier enabled state. Valid in GET operation. `dry_run_paused` Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. Valid in GET operation. 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.
surge_as_normal	boolean	Indicates whether or not to set the surge values as historical values. This field is no longer supported. Use update_baseline_from_surge instead.
surge_usage	surge_usage	Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.
suspect_files	array[suspect_files]	
typical_usage	typical_usage	Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

Name	Type	Description
update_baseline_from_surge	boolean	Sets the observed surge value as the new baseline on a volume.
workload	anti_ransomware_volume_workload	

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.

asynchronous_directory_delete

Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.

Name	Type	Description
enabled	boolean	Specifies whether asynchronous directory delete from the client is enabled on the volume.
trash_bin	string	Name of the trash bin directory. If no "trash_bin" property is specified when enabling, the default trash bin name, "._ontaptrashbin", is used.

autosize

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

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

snapshot_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

parent_svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM. This field cannot be specified in a PATCH method.
uuid	string	The unique identifier of the SVM. This field cannot be specified in a PATCH method.

parent_volume

Name	Type	Description
_links	_links	
name	string	The name of the volume. This field cannot be specified in a PATCH method.

Name	Type	Description
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6 • x-nullable: true

clone

Name	Type	Description
has_flexclone	boolean	Specifies whether this volume is a parent of any FlexClone volume.
inherited_physical_used	integer	Inherited physical used from the clone's base snapshot.
inherited_savings	integer	Inherited savings from the clone's base snapshot.
is_flexclone	boolean	Specifies if this volume is a normal FlexVol volume or FlexClone volume. This field needs to be set when creating a FlexClone volume. Valid in POST.
parent_snapshot	snapshot_reference	
parent_svm	parent_svm	SVM, applies only to SVM-scoped objects.
parent_volume	parent_volume	
split_complete_percent	integer	Percentage of FlexClone volume blocks split from its parent volume.
split_estimate	integer	Space required by the containing-aggregate to split the FlexClone volume.

Name	Type	Description
split_initiated	boolean	This field is set when a split is executed on any FlexClone volume, that is when the FlexClone volume is split from its parent FlexVol volume. Setting this field initiates a split of a FlexClone volume from a FlexVol volume. 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 volume 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 volume constituents is 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.
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 autodelete is currently enabled on this volume.
reserve_percent	integer	The space that has been set aside as a reserve for snapshot usage, in percent.
used	integer	The total space used by snapshots 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.

Name	Type	Description
block_storage_inactive_user_data	integer	The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store.
capacity_tier_footprint	integer	Space used by capacity tier for this volume in the FabricPool aggregate, in bytes.
footprint	integer	Data used for this volume in the aggregate, in bytes.
large_size_enabled	boolean	Specifies whether the support for large volumes and large files is enabled on the volume.
local_tier_footprint	integer	Space used by the local tier for this volume in the aggregate, in bytes.
logical_space	logical_space	
max_size	string	Indicates the maximum size supported for each constituent of the FlexGroup volume.
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	

Name	Type	Description
total_footprint	integer	Data and metadata used for this volume in the aggregate, in bytes.
total_metadata	integer	Space used by the total metadata in the volume, in bytes.
total_metadata_footprint	integer	Space used by the volume metadata footprint 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	
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 volume constituents is not supported.
name	string	FlexGroup volume constituent name.
space	space	

idcs_scanner

Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner.

Name	Type	Description
enabled	boolean	Specifies the administrative state of the inactive data compression scanner. Disabling inactive data compression is not allowed on Capacity optimized Flash with QAT supported platforms.
inactive_days	integer	Data blocks older than, or equal to, 'inactive_days' are picked up by the inactive data compression scanner. Valid for PATCH only. Only applicable when 'operation_state' set to 'active'.
mode	string	Specifies the mode of inactive data compression scanner. Valid for PATCH and GET.
operation_state	string	Specifies the operational state of the inactive data compression scanner. VALID for PATCH and GET. Valid options for PATCH are "idle" and "active".
status	string	Status of last inactive data compression scan on the volume.
threshold_inactive_time	string	Time interval after which inactive data compression is automatically triggered. The value is in days and is represented in the ISO-8601 format "P<num>D", for example "P3D" represents a duration of 3 days. This field is not supported on QAT supported platforms.</num>

policy

Name	Type	Description
name	string	Specifies the name of the efficiency policy. The "inline-only" and "none" policies are not supported on Capacity optimized Flash with QAT supported platform.

scanner

Name	Type	Description
compression	boolean	Start compression if scanning old data. Valid for PATCH and GET. This option is not supported for FSX/CVO platforms.
dedupe	boolean	Start deduplication if scanning old data. Valid for PATCH and GET.
scan_old_data	boolean	Indicates whether or not to scan old data. Valid for PATCH and GET.
state	string	State of the volume efficiency scanner. Valid for PATCH and GET. Valid options for PATCH are "idle" and "active".

space_savings

Name	Type	Description
compression	integer	Total disk space that is saved by compressing blocks on the referenced file system, in bytes.
compression_percent	integer	Percentage of total disk space that is saved by compressing blocks on the referenced file system.
dedupe	integer	Total disk space that is saved by deduplication and file cloning, in bytes.
dedupe_percent	integer	Percentage of total disk space that is saved by deduplication and file cloning.
dedupe_sharing	integer	Total disk space that is shared due to deduplication and file cloning.
total	integer	Total disk space saved in the volume due to deduplication, compression and file cloning, in bytes.

Name	Type	Description
total_percent	integer	Percentage of total disk space saved in the volume due to deduplication, compression and file cloning.

efficiency

Name	Type	Description
application_io_size	string	Block size to use by compression. 8k and auto are only allowed for POST. Only auto is supported on QAT supported platforms.
auto_state	string	Automatic deduplication schedule volume state. auto ‐ Volumes with auto_state set to auto start post-process deduplication automatically. deprioritized ‐ Volumes with auto_state set to deprioritized do not start post-process deduplication automatically.
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 FlexGroup volumes, where some of the constituent volumes are compaction enabled and some are disabled.

Name	Type	Description
compression	string	The system can be enabled/disabled compression. Disabling compression is not allowed on Capacity optimized Flash with QAT supported platforms. inline ‐ Data will be compressed first and written to the volume. background ‐ Data will be written to the volume and compressed later. both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run. none ‐ None mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compression enabled and some are disabled. that On volumes with container compression enabled, background compression refers to inactive data compression scan enabled on the volume.
compression_type	string	Compression type to use by compression. Valid for PATCH and GET.

Name	Type	Description
cross_volume_dedupe	string	The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. Disabling cross volume dedupe is not allowed on Capacity optimized Flash with QAT supported platforms. inline ‐ Data will be cross volume deduped first and written to the volume. background ‐ Data will be written to the volume and cross volume deduped later. both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.
dedupe	string	The system can be enabled/disabled dedupe. Disabling dedupe is not allowed on Capacity optimized Flash with QAT supported platforms. 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 FlexGroup volumes, where some of the constituent volumes are dedupe enabled and some are disabled.
has_savings	boolean	When true, indicates that the volume contains shared(deduplication, file clones) or compressed data.

Name	Type	Description
idcs_scanner	idcs_scanner	Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner.
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.
logging_enabled	boolean	When true, indicates that space savings for any newly-written data are being logged.
op_state	string	Sis status of the volume.
policy	policy	
progress	string	Sis progress of the volume.
scanner	scanner	
schedule	string	Schedule associated with volume.
space_savings	space_savings	

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 and Cloud Volumes 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 and Cloud Volumes ONTAP ‐ enabled ‐ All supported storage efficiency features for the volume are enabled. custom ‐ Read-only field currently only supported for the FSx for ONTAP and Cloud Volumes ONTAP, user-defined storage efficiency features are enabled. For other platforms ‐ enabled ‐ At least one storage efficiency feature for the volume is enabled. • enum: ["disabled", "enabled", "mixed", "custom"] • Introduced in: 9.9 • x-nullable: true
storage_efficiency_mode	string	Storage efficiency mode used by volume. This parameter is supported only on AFF platforms. There is no difference between default and efficient modes on QAT supported platforms and auto adaptive compression is set irrespective of the modes.
type	string	Sis Type of the volume.
volume_path	string	Absolute volume path of the volume.

status

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

encryption

Name	Type	Description
action	string	This field can be used to pause an ongoing rekey or conversion operation or resume a paused rekey or conversion operation. Valid in PATCH. 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.
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_create_time	string	Encryption key creation time of the volume.

Name	Type	Description
key_id	string	The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key.
rekey	boolean	If set to 'true', re-encrypts the volume with a new key. Valid in PATCH.
state	string	Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one.
status	status	
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.

Name	Type	Description
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.

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

throughput

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

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

metric

Performance numbers, such as IOPS, latency and throughput.

Name	Type	Description
_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 volume 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 dight), first digit is assumed to be zero.

policy

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

Name	Type	Description
_links	_links	
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

Name	Type	Description
min_throughput_iops	integer	Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_mbps	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
name	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
uuid	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

qos

QoS information

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

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume 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 FlexVol volumes or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVol volumes or FlexGroup volumes when the quota state is "off".
state	string	Quota state of the volume

last_error

Error information for the last failed file move on the constituent.

Name	Type	Description
code	integer	Error code of the last file move error on the constituent.
destination	integer	DSID of the destination constituent of the last file move error on the constituent.
file_id	integer	File ID of the last file move error on the constituent.
time	string	Time of the last file move error on the constituent.

movement

Properties on this constituent related to file movement.

Name	Type	Description
file_moves_started	integer	Number of file moves started on this constituent.
last_error	last_error	Error information for the last failed file move on the constituent.

Name	Type	Description
most_recent_start_time	string	Start time of the most recent file move on the constituent.

blocks_skipped

Number of blocks skipped by the scanner on this constituent due to various reasons.

Name	Type	Description
efficiency_blocks	integer	Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.
efficiency_percent	integer	Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.
fast_truncate	integer	Number of blocks skipped by the scanner on this constituent because fast truncate is currently running on files.
footprint_invalid	integer	Number of blocks skipped by the scanner on this constituent because of files with invalid space footprints.
in_snapshot	integer	Number of blocks skipped by the scanner on this constituent because of files in snapshots.
incompatible	integer	Number of blocks skipped by the scanner on this constituent because of incompatible files.
metadata	integer	Number of blocks skipped by the scanner on this constituent because of metadata files.
on_demand_destination	integer	Number of blocks skipped by the scanner on this constituent because of on demand destination files.

Name	Type	Description
other	integer	Number of blocks skipped by the scanner on this constituent for all other reasons.
remote_cache	integer	Number of blocks skipped by the scanner on this constituent because of remote caches.
too_large	integer	Number of blocks skipped by the scanner on this constituent because of files that are larger than rebalancing.max_file_size.
too_small	integer	Number of blocks skipped by the scanner on this constituent because of files that are smaller than rebalancing.min_file_size.
write_fenced	integer	Number of blocks skipped by the scanner on this constituent because of files fenced for write operations.

files_skipped

Number of files skipped by the scanner on this constituent due to various reasons.

Name	Type	Description
efficiency_blocks	integer	Number of files skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.
efficiency_percent	integer	Number of files skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.
fast_truncate	integer	Number of files skipped by the scanner on this constituent because fast truncate is currently running on the file.
footprint_invalid	integer	Number of files skipped by the scanner on this constituent because their space footprints are invalid.

Name	Type	Description
in_snapshot	integer	Number of files skipped by the scanner on this constituent because they are trapped in snapshots.
incompatible	integer	Number of files skipped by the scanner on this constituent because they are incompatible.
metadata	integer	Number of files skipped by the scanner on this constituent because they metadata files.
on_demand_destination	integer	Number of files skipped by the scanner on this constituent because they are on demand destinations.
other	integer	Number of files skipped by the scanner on this constituent for all other reasons.
remote_cache	integer	Number of files skipped by the scanner on this constituent because they are remote caches.
too_large	integer	Number of files skipped by the scanner on this constituent because they are larger than rebalancing.max_file_size.
too_small	integer	Number of files skipped by the scanner on this constituent because they are smaller than rebalancing.min_file_size.
write_fenced	integer	Number of files skipped by the scanner on this constituent because they are fenced for write operations.

scanner

Properties related to determining which files to move and where to move them to.

Name	Type	Description
blocks_scanned	integer	Number of blocks scanned on this constituent.
blocks_skipped	blocks_skipped	Number of blocks skipped by the scanner on this constituent due to various reasons.
files_scanned	integer	Number of files scanned on this constituent.
files_skipped	files_skipped	Number of files skipped by the scanner on this constituent due to various reasons.

engine

Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions.

Name	Type	Description
movement	movement	Properties on this constituent related to file movement.
scanner	scanner	Properties related to determining which files to move and where to move them to.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

rebalancing

Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

Name	Type	Description
_links	_links	
data_moved	integer	The amount of data that has been moved in or out of a constituent. A positive value represents data moving into the constituent while a negative value is data moving out of the constituent.
engine	engine	Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions.
exclude_snapshots	boolean	Specifies whether or not to exclude files that are stuck in snapshots during rebalancing operation. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "exclude_snapshots" value. Once the operation is started, any changes to the "exclude_snapshots" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "exclude_snapshots" value.
imbalance_percent	integer	Represents the percentage the volume is out of balance.
imbalance_size	integer	Represents how much the volume is out of balance, in bytes.

Name	Type	Description
max_constituent_imbalance_percent	integer	Absolute percentage of the constituent that is most out of balance. This value will update every 30 seconds when rebalancing is not active and every 10 seconds when rebalancing is active.
max_file_moves	integer	Specifies the maximum number of concurrent file moves in a volume capacity rebalancing operation on a constituent of the FlexGroup volume. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "max_file_moves" value. Once the operation is started, any changes to the "max_file_moves" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "max_file_moves" value.
max_runtime	string	This optional field specifies the maximum time a capacity rebalancing operation runs for. Once the maximum runtime has passed, the capacity rebalancing operation stops. If it is not set, the default value is 6 hours. This value cannot be updated while a capacity rebalancing operation is running. The maximum runtime 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 "P3D" represents a duration of 3 days. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively.</num></num></num></num></num>

Name	Type	Description
max_threshold	integer	Specifies the maximum imbalance percentage for FlexGroup volume constituents. When a constituent's imbalance percentage is larger than this value, files are moved from the constituent. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "max_threshold" value. Once the operation is started, any changes to the "max_threshold" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "max_threshold" value.
min_file_size	integer	Specifies the minimum file size to consider for a volume capacity rebalancing operation. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "min_file_size" value. Once the operation is started, any changes to the "min_file_size" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "min_file_size" value. The value must be a multiple of 4KB. If it is not set, the default value is 100MB. Setting "min-file-size" to less than the default value leads to more files being moved. Moved files use granular data, which may impact read/write I/O performance.

Name	Type	Description
min_threshold	integer	Specifies the minimum imbalance percentage for FlexGroup volume constituents. When a constituent's imbalance percentage is smaller than this value, files are not moved from the constituent. When a new capacity rebalancing operation is started on a FlexGroup volume, it will use the current "min_threshold" value. Once the operation is started, any changes to the "min_threshold" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "min_threshold" value.
notices	array[error]	Capacity rebalancing notice messages.
runtime	string	Duration the capacity rebalancing operation has been running.
start_time	string	Time when the current capacity rebalancing operation started, or when a future scheduled rebalancing operation begins.

Name	Type	Description
state	string	State of the volume capacity rebalancing operation. PATCH the state to "starting" to trigger the capacity rebalance operation, and include start_time to schedule rebalancing. PATCH the state to "stopping" to stop the capacity rebalance operation, or cancel a scheduled rebalancing operation. PATCH without the state with a valid start_time to modify the start_time of an existing scheduled rebalance operation. While a FlexGroup volume is rebalancing, every constituent will have a rebalancing engine that can either be scanning the filesystem for space usage and files to move, actively moving files or temporarily doing neither. If one or more constituents has a state of "rebalancing_source" or "rebalancing_dest", then files are being moved to rebalance the FlexGroup volume. If no files are being moved, more information about what the rebalancing engine is doing for each constituent is available using the "rebalancing.engine" property. The following values apply to FlexGroup volumes. not_running ‐ capacity rebalancing is not running on the volume. starting ‐ used in a PATCH operation to start a capacity rebalancing operation. rebalancing ‐ capacity rebalancing is running on the volume. paused ‐ volume capacity rebalancing is paused on the volume. stopping ‐ used in a PATCH operation to stop a capacity rebalancing operation. unknown ‐ the system was unable to determine the rebalancing state for the volume. The following values apply to FlexGroup volume constituents.

Name	Type	Description
stop_time	string	Time when the capacity rebalancing operation stopped.
target_used	integer	Represents the ideal used size of each constituent. Calculated by dividing the total FlexGroup volume used size by the number of constituents.
used_for_imbalance	integer	Represents the used size of each constituent, as determined by the rebalancing engine. Calculated by subtracting the size used by snapshots, the size of files pending deletion and the size of filesystem metadata from the volume used size.

retention

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

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

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

snaplock

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

Name	Type	Description
autocommit_period	string	Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num>
compliance_clock_time	string	This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
expiry_time	string	Expiry time of the volume. For volumes with an infinite, indefinite or unspecified SnapLock expiry time, "9999-12-31T23:59:59" is used to denote the time.
is_audit_log	boolean	Indicates if this volume has been configured as SnapLock audit log volume for the SVM .

Name	Type	Description
litigation_count	integer	Litigation count indicates the number of active legal-holds on the volume.
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.

Name	Type	Description
is_ontap	boolean	Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to an ONTAP destination. • readOnly: 1 • Introduced in: 9.9 • x-nullable: true

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 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.

Name	Type	Description
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.
used_by_snapshots	integer	Size that is logically used across all snapshots 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.

autodelete

Name	Type	Description
commitment	string	By default, snapshot autodelete does not delete snapshots locked by Snapmirror, clones of a volume, a LUN, an NVMe namespace, or a file. Deletion of snapshots locked by these applications is specified using this option. The default value is try.
defer_delete	string	Allows the user to inform snapshot autodelete to defer the deletion of a specified snapshot until the end. The default value is user_created.

Name	Type	Description
delete_order	string	Specifies the order in which snapshot autodelete occurs. Ordering is done using the date and time the snapshot is created. The default value is oldest_first.
enabled	boolean	Specifies whether snapshot autodelete is currently enabled on this volume.
prefix	string	Specifies the prefix of the snapshot which if matched, is deleted last. Used with autodelete_defer_delete when used with a prefix value.
target_free_space	integer	Snapshots are deleted, one at a time, until the used volume space reaches the value specified. The default is 20% free space or 80% utilized.
trigger	string	Specifies when the system should trigger an autodelete of snapshots. When set to *volume*, autodelete is triggered based on volume fullness. When set to *snap_reserve*, autodelete is triggered based on snapshot reserve fullness. The default value is *volume*.

snapshot

Name	Type	Description
autodelete	autodelete	
autodelete_enabled	boolean	Specifies whether snapshot autodelete is currently enabled on this volume. This field will no longer be supported in a future release. Use autodelete.enabled instead.

Name	Type	Description
autodelete_trigger	string	Specifies when the system should trigger an autodelete of snapshots. When set to *volume*, autodelete is triggered based on volume fullness. When set to *snap_reserve*, autodelete is triggered based on snapshot reserve fullness. The default value is *volume*. This field will no longer be supported in a future release. Use autodelete.trigger instead.
reserve_available	integer	Size available for snapshots within the snapshot reserve, in bytes.
reserve_percent	integer	The space that has been set aside as a reserve for snapshot usage, in percent.
reserve_size	integer	Size in the volume that has been set aside as a reserve for snapshot usage, in bytes.
space_used_percent	integer	Percentage of snapshot reserve size that has been used.
used	integer	The total space used by snapshots in the volume, in bytes.

space

Name	Type	Description
afs_total	integer	Total size of AFS, excluding snap-reserve, in bytes.
auto_adaptive_compression_footprint_data_reduction	integer	Savings achieved due to Auto Adaptive Compression, in bytes.
available	integer	The available space, in bytes.
available_percent	integer	The space available, as a percent.

Name	Type	Description
block_storage_inactive_user_data	integer	The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store.
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.
capacity_tier_footprint_data_reduction	integer	Savings achieved in the space used by the capacity tier for this volume in the FabricPool aggregate, in bytes.
compaction_footprint_data_reduction	integer	Savings achieved due to Data Compaction, 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.

Name	Type	Description
effective_total_footprint	integer	Volume footprint after efficiency savings, in bytes. effective total footprint represents total footprint after deducting auto adaptive compression and compaction savings. effective-footprint includes aggregate metadata used by volume.
expected_available	integer	Size that should be available for the volume, irrespective of available size in the aggregate, in bytes.
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.
is_used_stale	boolean	Specifies if the virtual space used is stale.
large_size_enabled	boolean	Indicates if the support for large FlexVol volumes and large files is enabled on this volume. When configured to true, FlexVol volume size can reach up to 300TB and single file size can reach 128TB.

Name	Type	Description
local_tier_footprint	integer	Space used by the local tier for this volume in the aggregate, in bytes.
logical_space	logical_space	
max_size	string	Indicates the maximum size supported for the FlexVol volume or for each constituent of the FlexGroup volume.
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.
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 is effective total footprint which is equivalent to total footprint after deducting auto adaptive compression and compaction savings. Physical-used also includes aggregate metadata used by volume.
physical_used_percent	integer	Size that is physically used in the volume, as a percentage.

Name	Type	Description
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 snapshots 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 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.
total_metadata	integer	Space used by the total metadata in the volume, in bytes.
total_metadata_footprint	integer	Space used by the volume metadata footprint 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.
volume_guarantee_footprint	integer	Space reserved for future writes in the volume, in bytes.

access

Raw count and latency data for access operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

audit

Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

dir

Raw count and latency data for directory-create operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

file

Raw count and latency data for file-create operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

other

Raw count and latency data for create operations on objects other than files, directories and symlinks.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

symlink

Raw count and latency data for symlink-create operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

create

Raw count and latency data for create operations.

Name	Type	Description
dir	dir	Raw count and latency data for directory-create operations.
file	file	Raw count and latency data for file-create operations.
other	other	Raw count and latency data for create operations on objects other than files, directories and symlinks.
symlink	symlink	Raw count and latency data for symlink-create operations.

getattr

Raw count and latency data for getattr operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

link

Raw count and latency data for link operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

lock

Raw count and latency data for lock operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

lookup

Raw count and latency data for lookup operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

open

Raw count and latency data for open operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

read

Raw count and latency data for read operations, including histograms categorizing operations by size and latency.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • readOnly: 1 • x-ntap-advanced: true • Introduced in: 9.11 • x-nullable: true
volume_protocol_latency_histogram_counts	array[integer]	
volume_protocol_latency_histogram_labels	array[string]	Labels for the latency histogram, ranging from <2us to >20s.
volume_protocol_size_histogram_counts	array[integer]	
volume_protocol_size_histogram_labels	array[string]	Labels for the size histogram, ranging from <4KB to >1024KB.

readdir

Raw count and latency data for readdir operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

readlink

Raw count and latency data for readlink operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

rename

Raw count and latency data for rename operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

setattr

Raw count and latency data for setattr operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

unlink

Raw count and latency data for unlink operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

watch

Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

write

Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • readOnly: 1 • x-ntap-advanced: true • Introduced in: 9.11 • x-nullable: true
volume_protocol_latency_histogram_counts	array[integer]	
volume_protocol_latency_histogram_labels	array[string]	Labels for the latency histogram, ranging from <2us to >20s.
volume_protocol_size_histogram_counts	array[integer]	
volume_protocol_size_histogram_labels	array[string]	Labels for the size histogram, ranging from <4KB to >1024KB.

cifs_ops_raw

Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

Name	Type	Description
access	access	Raw count and latency data for access operations.
audit	audit	Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.
create	create	Raw count and latency data for create operations.
getattr	getattr	Raw count and latency data for getattr operations.
link	link	Raw count and latency data for link operations.
lock	lock	Raw count and latency data for lock operations.
lookup	lookup	Raw count and latency data for lookup operations.
open	open	Raw count and latency data for open operations.
read	read	Raw count and latency data for read operations, including histograms categorizing operations by size and latency.
readdir	readdir	Raw count and latency data for readdir operations.
readlink	readlink	Raw count and latency data for readlink operations.
rename	rename	Raw count and latency data for rename operations.
setattr	setattr	Raw count and latency data for setattr operations.
unlink	unlink	Raw count and latency data for unlink operations.

Name	Type	Description
watch	watch	Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.
write	write	Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

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.

Name	Type	Description
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.

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

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.

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

nfs_ops_raw

Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

Name	Type	Description
access	access	Raw count and latency data for access operations.
audit	audit	Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.
create	create	Raw count and latency data for create operations.
getattr	getattr	Raw count and latency data for getattr operations.

Name	Type	Description
link	link	Raw count and latency data for link operations.
lock	lock	Raw count and latency data for lock operations.
lookup	lookup	Raw count and latency data for lookup operations.
open	open	Raw count and latency data for open operations.
read	read	Raw count and latency data for read operations, including histograms categorizing operations by size and latency.
readdir	readdir	Raw count and latency data for readdir operations.
readlink	readlink	Raw count and latency data for readlink operations.
rename	rename	Raw count and latency data for rename operations.
setattr	setattr	Raw count and latency data for setattr operations.
unlink	unlink	Raw count and latency data for unlink operations.
watch	watch	Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.
write	write	Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

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
cifs_ops_raw	cifs_ops_raw	Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.
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.

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.
nfs_ops_raw	nfs_ops_raw	Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_ delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

Name	Type	Description
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
_links	_links	
name	string	The name of the SVM. This field cannot be specified in a PATCH method.
uuid	string	The unique identifier of the SVM. This field cannot be specified in a PATCH method.

tiering

Name	Type	Description
min_cooling_days	integer	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.
object_tags	array[string]	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".

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

returned_error

Name	Type	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. The PATCH API can also be used to start or stop non-disruptive volume capacity rebalancing for FlexGroup volumes in addition to modifying capacity rebalancing properties. 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.

Platform Specifics

Unified ONTAP: PATCH must be used to update the attributes of a volume.

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.
- <code>encryption.action</code> - 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 conversion pause`

- `volume encryption conversion resume`
- `volume encryption rekey start`
- `volume encryption rekey pause`
- `volume encryption rekey resume`
- `volume rebalance start`
- `volume rebalance stop`
- `volume rebalance modify`
- `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`
- `volume file async-delete client disable`
- `volume file async-delete client enable`

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 to restore volume to the point in time the snapshot was taken.
restore_to.snapshot.name	string	query	False	Name of the snapshot to restore volume to the point in time the snapshot was taken.
restore_to.path	string	query	False	Path to the file which is restored from the snapshot. • Introduced in: 9.10

Name	Type	In	Required	Description
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
restore_to.byte_count	number	query	False	Number of bytes to restore from the source file, in multiples of 4096. • Introduced in: 9.10
preserve_lun_ids	boolean	query	False	Specifies whether LUN IDs need to be preserved during a snapshot restore operation. • Introduced in: 9.10 • Default value:

Name	Type	In	Required	Description
nvfail	string	query	False	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. • Introduced in: 9.10 • enum: ["off", "on"]
snapshot_directory_access_enabled	boolean	query	False	This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume. • Introduced in: 9.10 • Default value: 1

Name	Type	In	Required	Description
sizing_method	string	query	False	Represents the method to modify the size of a Flexgroup. The following methods are supported: • use_existing_resources - Increases or decreases the size of the FlexGroup by increasing or decreasing the size of the current FlexGroup resources • add_new_resources - Increases the size of the FlexGroup by adding new resources. 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	Naming Scheme for automatic snapshots: • create_time - Automatic snapshots 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"] </scheduled_frequency>
validate_only	boolean	query	False	Validate the volume move or volume conversion operations and their parameters, without actually performing the operation. • Introduced in: 9.13

Name	Type	In	Required	Description
clone.match_parent_storage_tier	boolean	query	False	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. • Introduced in: 9.9

Name	Type	In	Required	Description
aggressive_readahead_mode	string	query	False	Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. This feature is only available on FabricPool volumes on FSx for ONTAP and Cloud Volumes ONTAP. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress. • Introduced in: 9.13 • Default value: 1 • enum: ["none", "file_prefetch", "sequential_read", "cross_file_sequential_read"]

Name	Type	In	Required	Description
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202. • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
_tags	array[string]	Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings.
access_time_enabled	boolean	Indicates whether or not access time updates are enabled on the volume.
activity_tracking	activity_tracking	

Name	Type	Description
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
aggressive_readahead_mode	string	Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress.
analytics	analytics	
anti_ransomware	anti_ransomware	Anti-ransomware related information of the volume.
application	application	
asynchronous_directory_delete	asynchronous_directory_delete	Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.
autosize	autosize	
clone	clone	

Name	Type	Description
cloud_retrieval_policy	string	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.
cloud_write_enabled	boolean	Indicates whether or not cloud writes are enabled on the volume. NFS writes to this volume are sent to the cloud directly instead of the local performance tier. This feature is only available on volumes in FabricPools.
comment	string	A comment for the volume. Valid in POST or PATCH.
consistency_group	consistency_group	Consistency group the volume is part of.

Name	Type	Description
constituents	array[constituents]	FlexGroup volume constituents. FlexGroup volume constituents can be retrieved more efficiently by specifying "is_constituent=true" or "is_constituent=true&flexgroup.uuid=<flexgroup.uuid>" as query parameters.</flexgroup.uuid>
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. If a volume is being created on multiple aggregates and the "constituents_per_aggregate" field is not specified, the default value of the "constituents_per_aggregate" field is 4. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list, unless 'optimize_aggregates' is specified as 'true'. If the "aggregates.name" or "aggregates.uuid" is specified in a PATCH request to expand an existing FlexGroup volume, the default value of the "constituents_per_aggregate" field is 1. The volume expand operation is only supported on FlexGroup volumes.
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.

Name	Type	Description
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	
granular_data	boolean	State of granular data on the volume. This setting is `true` by default when creating an AWS S3 FlexGroup volume via a POST request and `false` by default for creating any other type of FlexGroup volume via a POST request. On FlexVol volumes, the setting is always `false`, as only FlexGroup volumes support this feature. Once enabled, this setting can only be disabled by restoring a snapshot. Earlier versions of ONTAP (pre 9.12) are not compatible with this feature. Therefore, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the setting was enabled. • Introduced in: 9.12 • x-nullable: true

Name	Type	Description
granular_data_mode	string	Mode of granular data on the volume. This setting defaults to `basic` when the `granular_data` parameter is set to `true`, but can be specified at the time of creation via a POST request. Earlier versions of ONTAP (pre 9.12) are not compatible with the `basic` setting. Therefore, when set to `basic`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `basic` mode was enabled. Earlier versions of ONTAP (pre 9.16) are not compatible with the `advanced` setting. Therefore, when set to `advanced`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `advanced` mode was enabled. • enum: ["disabled", "basic", "advanced"] • Introduced in: 9.16 • x-nullable: true
guarantee	guarantee	
is_object_store	boolean	Specifies whether the volume is provisioned for an object store server.
is_svm_root	boolean	Specifies whether the volume is a root volume of the SVM it belongs to.
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.

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 volume constituents are not supported.
msid	integer	The volume's Master Set ID.
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 FlexGroup volumes, 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.
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume.
rebalancing	rebalancing	Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

Name	Type	Description
scheduled_snapshot_naming_scheme	string	Naming Scheme for automatic snapshots: • create_time - Automatic snapshots 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 snapshots in the volume.
snapshot_directory_access_enabled	boolean	This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume.
snapshot_locking_enabled	boolean	Specifies whether or not snapshot copy locking is enabled on the volume.
snapshot_policy	snapshot_policy	This is a reference to the snapshot 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.
status	array[string]	Describes the current status of a volume.
style	string	The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates or license. 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 volume. When the UDO License is installed, and no aggregates are specified, the system automatically provisions a FlexGroup volume on system selected aggregates. 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 ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroup volumes flexgroup_constituent ‐ FlexGroup volume constituents.

Name	Type	Description
tiering	tiering	
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6 • x-nullable: true
validate_only	boolean	Validate the volume move or volume conversion operations and their parameters, without actually performing the operation.

Example request


```
{
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "activity_tracking": {
    "state": "string",
    "unsupported_reason": {
      "code": "124518405",
      "message": "Volume activity tracking cannot be enabled on volumes
that contain LUNs."
    }
  },
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "aggressive_readahead_mode": "string",
  "analytics": {
    "files_scanned": 43002,
    "initialization": {
      "state": "string"
    },
    "scan_progress": 17,
    "scan_throttle_reason": {
      "arguments": [
        "string"
      ],
      "code": "6739881",
      "message": "The file system analytics scan running on volume
\"fsavol2\" in SVM \"vs2\" has slowed down on node \"my_node\". Reason:
Computing resources are being used by higher priority workloads."
    },
    "state": "string",
    "total_files": 101890,
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
    }
  },
  "anti_ransomware": {
```

```
    "attack_detection_parameters": {
      "file_create_op_rate_surge_notify_percent": 100
    },
    "attack_probability": "string",
    "attack_reports": [
      {
        "time": "2021-06-01 11:06:41 -0400"
      }
    ],
    "dry_run_start_time": "string",
    "space": {
      "snapshot_count": 0,
      "used": 0,
      "used_by_logs": 0,
      "used_by_snapshots": 0
    },
    "state": "string",
    "surge_usage": {
      "file_create_peak_rate_per_minute": 10,
      "file_delete_peak_rate_per_minute": 50,
      "file_rename_peak_rate_per_minute": 30,
      "high_entropy_data_write_peak_percent": 30,
      "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
      "time": "2021-12-01 12:46:20 -0500"
    },
    "suspect_files": [
      {
        "count": 0,
        "entropy": "string",
        "format": "string"
      }
    ],
    "typical_usage": {
      "file_create_peak_rate_per_minute": 50,
      "file_delete_peak_rate_per_minute": 10,
      "file_rename_peak_rate_per_minute": 5,
      "high_entropy_data_write_peak_percent": 10,
      "high_entropy_data_write_peak_rate_kb_per_minute": 1200
    },
    "workload": {
      "file_extension_types_count": 3,
      "file_extensions_observed": [
        "pdf",
        "jpeg",
        "txt"
      ],
```

```
      "historical_statistics": {
        "file_create_peak_rate_per_minute": 50,
        "file_delete_peak_rate_per_minute": 10,
        "file_rename_peak_rate_per_minute": 5,
        "high_entropy_data_write_peak_percent": 10,
        "high_entropy_data_write_peak_rate_kb_per_minute": 1200
      },
      "newly_observed_file_extensions": [
        {
          "count": [
            20
          ],
          "name": [
            "lockile"
          ]
        }
      ],
      "surge_statistics": {
        "file_create_peak_rate_per_minute": 10,
        "file_delete_peak_rate_per_minute": 50,
        "file_rename_peak_rate_per_minute": 30,
        "high_entropy_data_write_peak_percent": 30,
        "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
        "time": "2021-12-01 12:46:20 -0500"
      },
      "surge_usage": {
        "file_create_peak_rate_per_minute": 10,
        "file_delete_peak_rate_per_minute": 50,
        "file_rename_peak_rate_per_minute": 30,
        "high_entropy_data_write_peak_percent": 30,
        "high_entropy_data_write_peak_rate_kb_per_minute": 2500,
        "newly_observed_file_extensions": [
          {
            "count": [
              20
            ],
            "name": [
              "lockile"
            ]
          }
        ],
        "time": "2021-12-01 12:46:20 -0500"
      },
      "typical_usage": {
        "file_create_peak_rate_per_minute": 50,
        "file_delete_peak_rate_per_minute": 10,
```

```
        "file_rename_peak_rate_per_minute": 5,
        "high_entropy_data_write_peak_percent": 10,
        "high_entropy_data_write_peak_rate_kb_per_minute": 1200
      }
    }
  },
  "application": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "asynchronous_directory_delete": {
    "trash_bin": "string"
  },
  "autosize": {
    "mode": "string"
  },
  "clone": {
    "inherited_physical_used": 0,
    "inherited_savings": 0,
    "is_flexclone": true,
    "split_complete_percent": 0,
    "split_estimate": 0
  },
  "cloud_retrieval_policy": "string",
  "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": {
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "percent_complete": 0,
        "state": "replicating",
        "tiering_policy": "string"
      },
      "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
        },
        "max_size": "string",
        "metadata": 0,
        "over_provisioned": 0,
        "performance_tier_footprint": 0,
        "snapshot": {
          "used": 0
        },
        "total_footprint": 0,
        "total_metadata": 0,
        "total_metadata_footprint": 0,
        "used": 0
      }
    }
  ],
  "create_time": "2018-06-04 15:00:00 -0400",
  "efficiency": {
    "application_io_size": "string",
    "auto_state": "string",
    "compaction": "string",
    "compression": "string",
    "compression_type": "string",
    "cross_volume_dedupe": "string",
    "dedupe": "string",
    "idcs_scanner": {
      "mode": "string",
      "operation_state": "string",
      "status": "string",
      "threshold_inactive_time": "P14D"
    },
    "last_op_begin": "string",
    "last_op_end": "string",
    "last_op_err": "string",
    "last_op_size": 0,
    "last_op_state": "string",
    "op_state": "string",
    "policy": {
```

```
      "name": "string"
    },
    "progress": "string",
    "scanner": {
      "state": "string"
    },
    "schedule": "string",
    "space_savings": {
      "compression": 0,
      "compression_percent": 0,
      "dedupe": 0,
      "dedupe_percent": 0,
      "dedupe_sharing": 0,
      "total": 0,
      "total_percent": 0
    },
    "state": "string",
    "storage_efficiency_mode": "string",
    "type": "string",
    "volume_path": "string"
  },
  "encryption": {
    "action": "string",
    "key_create_time": "2022-01-01 14:00:00 -0500",
    "key_id": "string",
    "state": "string",
    "status": {
      "code": "string",
      "message": "string"
    },
    "type": "string"
  },
  "files": {
    "used": 0
  },
  "flash_pool": {
    "cache_eligibility": "string",
    "cache_retention_priority": "string",
    "caching_policy": "string"
  },
  "flexcache_endpoint_type": "string",
  "flexgroup": {
    "name": "my_flexgroup",
    "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
  },
  "granular_data_mode": "string",
```

```
  "guarantee": {
    "type": "string"
  },
  "movement": {
    "capacity_tier_optimized": true,
    "cutover_window": 30,
    "destination_aggregate": {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": 0,
    "start_time": "2020-12-07 03:45:12 -0500",
    "state": "replicating",
    "tiering_policy": "string"
  },
  "name": "vol_cs_dept",
  "nas": {
    "export_policy": {
      "id": 100,
      "name": "default"
    },
    "junction_parent": {
      "name": "vs1_root",
      "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
    },
    "path": "/user/my_volume",
    "security_style": "string",
    "unix_permissions": 493
  },
  "optimize_aggregates": true,
  "qos": {
    "policy": {
      "max_throughput_iops": 10000,
      "max_throughput_mbps": 500,
      "min_throughput_iops": 2000,
      "min_throughput_mbps": 500,
      "name": "performance",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "quota": {
    "state": "string"
  },
  "rebalancing": {
    "data_moved": 0,
    "engine": {
```

```
      "movement": {
        "file_moves_started": 0,
        "last_error": {
          "code": 0,
          "destination": 0,
          "file_id": 0,
          "time": "2018-06-04 15:00:00 -0400"
        },
        "most_recent_start_time": "2018-06-04 15:00:00 -0400"
      },
      "scanner": {
        "blocks_scanned": 0,
        "blocks_skipped": {
          "efficiency_blocks": 0,
          "efficiency_percent": 0,
          "fast_truncate": 0,
          "footprint_invalid": 0,
          "in_snapshot": 0,
          "incompatible": 0,
          "metadata": 0,
          "on_demand_destination": 0,
          "other": 0,
          "remote_cache": 0,
          "too_large": 0,
          "too_small": 0,
          "write_fenced": 0
        },
        "files_scanned": 0,
        "files_skipped": {
          "efficiency_blocks": 0,
          "efficiency_percent": 0,
          "fast_truncate": 0,
          "footprint_invalid": 0,
          "in_snapshot": 0,
          "incompatible": 0,
          "metadata": 0,
          "on_demand_destination": 0,
          "other": 0,
          "remote_cache": 0,
          "too_large": 0,
          "too_small": 0,
          "write_fenced": 0
        }
      }
    },
    "imbalance_percent": 0,
```

```
    "imbalance_size": 0,
    "max_constituent_imbalance_percent": 0,
    "max_runtime": "string",
    "notices": [
      {
        "arguments": [
          {
            "code": "string",
            "message": "string"
          }
        ],
        "code": "4",
        "message": "entry doesn't exist"
      }
    ],
    "runtime": "string",
    "start_time": "string",
    "state": "rebalancing",
    "stop_time": "string",
    "target_used": 0,
    "used_for_imbalance": 0
  },
  "scheduled_snapshot_naming_scheme": "string",
  "snaplock": {
    "append_mode_enabled": "",
    "autocommit_period": "P30M",
    "compliance_clock_time": "2018-06-04 15:00:00 -0400",
    "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"
    },
    "unspecified_retention_file_count": 10
  },
  "snapshot_count": 0,
  "snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "space": {
    "auto_adaptive_compression_footprint_data_reduction": 0,
    "available": 0,
```

```
    "block_storage_inactive_user_data": 0,
    "block_storage_inactive_user_data_percent": 0,
    "capacity_tier_footprint": 0,
    "capacity_tier_footprint_data_reduction": 0,
    "compaction_footprint_data_reduction": 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,
    "effective_total_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
    },
    "max_size": "string",
    "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": {
        "commitment": "string",
        "defer_delete": "string",
        "delete_order": "string",
        "prefix": "string",
        "trigger": "string"
      },
      "autodelete_trigger": "string",
      "reserve_available": 0,
      "reserve_size": 0,
      "space_used_percent": 0,
      "used": 0
    },
    "snapshot_reserve_unusable": 0,
```

```
      "snapshot_spill": 0,
      "total_footprint": 0,
      "total_metadata": 0,
      "total_metadata_footprint": 0,
      "used": 0,
      "user_data": 0,
      "volume_guarantee_footprint": 0
    },
    "state": "string",
    "status": [
      "string"
    ],
    "style": "string",
    "tiering": {
      "object_tags": [
        "string"
      ],
      "policy": "string",
      "supported": true
    },
    "use_mirrored_aggregates": true,
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
```


Response

```
Status: 200, Ok
```

Name	Type	Description
job	job_link	

Example response


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

Response

```
Status: 202, Accepted
```

Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
262256	Tags cannot be modified for the volume because of the specified reason.
524511	Failed to restore snapshot "restore_to.snapshot.name".
524849	Attribute cannot be modified for the target volume because of the specified reason.
787141	The specified "aggregates.name" and "aggregates.uuid" refer to different aggregates.
787144	The volume is on an aggregate that is not part of a FabricPool.
917505	SVM not found.
917534	The value specified for field "-size" is too small. Update the field "-size" with the minimum size allowed and retry.
917597	Provisioning a volume on a root aggregate is not supported.
917617	Cannot modify snapshot policy on volumes of type DC, DP, LS or TMP.
917628	Modification of the specified fields is not allowed for the target volume type.
917829	Volume autosize grow threshold must be larger than autosize shrink threshold.
917831	Volume minimum autosize must be smaller than the maximum autosize.
917888	Invalid character in volume name. It can have '_' and alphanumeric characters.;
918183	The specified tiering policy is not supported for the current system configuration.
918193	Cannot modify tiering min cooling days when vol move is in progress.

Error Code	Description
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 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.
918334	Cache cloud retrieval policy requires an effective cluster version of 9.14 or later.
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.
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.

Error Code	Description
918538	The match-parent-storage-tier option is not supported for clone creation.
918577	Only FlexVol to FlexGroup conversions are currently supported.
918583	Disabling large volume size is not supported because the volume contains files larger than 16TB or the volume size is larger than 100TB.
918599	Cloud write requires an effective cluster version of 9.13.1 or later.
918600	Cloud write is only supported for volumes in a FabricPool on FSx for ONTAP or Cloud Volumes ONTAP.
918601	Only the \"all\" tiering policy is allowed on a volume with cloud write.
918602	Cannot move volume when cloud write is enabled.
918603	Cloud write is only supported on volumes with Read-Write access.
918604	Cloud write is not supported because the volume is the destination or source endpoint of one or more SnapMirror relationships.
918605	Cloud write cannot be enabled on a clone volume.
918606	Cloud write cannot be enabled on a volume containing LUNs.
918607	Cloud write is not supported because this volume belongs to an SVM that is the source of a DR relationship.
918624	Tiering storage class requires an effective cluster version of ONTAP 9.13.1 or later.
918625	Tiering storage class is only supported for volumes in a FabricPool on FSx for ONTAP or Cloud Volumes ONTAP for AWS.
918626	Tiering storage class is not supported for volumes in a FlexGroup.
918627	The specified value for "tiering.storage_class" is not supported.
918636	Volume "name" in SVM "svm.name" is part of a Consistency Group. Snapshot locking cannot be enabled on a volume that is part of a Consistency Group.

Error Code	Description
918639	One or more of the fields "-max-size","-max-autosize","-size" are being set or modified. Set the values so that "-max-size" is always greater than or equal to "-max-autosize" and "-size".
918641	The value passed for "-max-size" is invalid. Valid values for the current platform are specified.
918647	One or more of the fields "-max-autosize","-size" are being set or modified. Set the values so that "-max-autosize" and "-size" are always less than or equal to the specified value.
918650	The asynchronous directory delete trash bin name cannot be specified unless asynchronous directory delete from the client is being enabled.
918652	"error" is an invalid value for field "-state". Valid values are "online", "offline" and "restricted".
918655	Cloud write is not supported on the current configuration.
918658	Volume is on a FabricPool-enabled aggregate. CIFS applications might time out when cloud write is enabled on a volume that is on a FabricPool-enabled aggregate.
918659	Cloud write can only be enabled on FabricPool-enabled aggregates.
918660	Cloud write enabled volumes can only be moved to FabricPool-enabled aggregates. Disable cloud write to move the volume to a non-FabricPool aggregates.
918703	Modification of the specified volume properties are not supported on this platform.
1638480	Failed to promote snapshot "restore_to.snapshot.name" because one or more newer snapshots are currently used as a reference snapshot for data protection operations.
1638590	Promoting a "-pre-conversion" snapshot is not supported.
2424998	Unable to determine whether MetroCluster is configured.
2621572	Modification of the specified field is not permitted on an SVM that is configured as the destination for identity discard SVM DR.
6488103	Async directory delete from the client is already enabled on volume "name" in SVM "svm.name".
6488104	Async directory delete from the client is already disabled on volume "name" in SVM "svm.name".

Error Code	Description
6881612	Cannot enable inline deduplication on a non AFF Platform or on non HYA aggregate.
6881640	The "efficiency.state" field is only supported on read-write volumes on FSx for ONTAP and Cloud Volumes ONTAP.
6881651	File compression is not supported on volumes with **max-size › 300TB**. Modify the volume max-size, then retry.
6881659	Changing "storage-efficiency-mode" from "efficient" to "default" is not supported.
6881661	Volume efficiency scanner cannot be started on volume with efficiency policy set as inline-only.
8454210	Failed to complete the operation because a QoS workload could not be created, modified or deleted. Wait a few seconds, and then try the operation again.
9437885	The volume is not online.
9437886	Volume snapshot restore failed because there are newer unexpired locked snapshots in the volume.
13107256	Operation is only supported on FlexGroup volumes.
13107349	Operation is only supported on flexible volumes and FlexGroup volumes.
13107371	Operation is only supported on read-write FlexGroup volumes.
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.
13107415	Failed to lookup a volume property.
13107431	Failed to lookup an SVM property.
13107433	A snapshot is scheduled to be taken within the volume capacity rebalancing runtime.
13107434	A SnapMirror update is scheduled within the volume capacity rebalancing runtime.
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.

Error Code	Description
13109258	Cannot enable granular data on volume "name" in Vserver "svm.name". This setting can only be enabled on FlexGroups.
13109259	Granular data cannot be disabled on volume "name" in Vserver "svm.name". This property can only be disabled by restoring a snapshot.
13109260	Failed to enable granular data on the volume.
13763477	Cannot modify the SnapLock configuration on volume "name" in SVM "svm.name" because the volume is a part of a consistency group.
65536965	The key manager on data Vserver \"{0}\" is in the blocked state due to key access errors. Possible reasons for a blocked state include the top-level external key protection key is disabled or not found or the current user does not have access to perform the requested action. Fix the issues with the key manager that manages the top-level external key protection key on data Vserver \"{0}\" to bring the key manager to the active state and then try to bring the volume online again.
65539432	Cannot online volume \"{1}\" on Vserver \"{0}\" while the keystore is being initialized. Wait until the keystore is in the active state, and rerun the volume operation.
65539433	Cannot online volume \"{1}\" on Vserver \"{0}\" while the keystore is being disabled.
66846758	Modification of the specified fields is not allowed for FlexCache volumes.
66846772	Modification of the specified fields is not allowed for origin of FlexCache volumes.
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.
111411207	Volume file system analytics is not supported on volumes that contain LUNs.
111411209	Volume file system analytics is not supported on FlexCache volumes.

Error Code	Description
111411210	Volume file system analytics is not supported on audit staging volumes.
111411211	Volume file system analytics is not supported on object store server volumes.
111411212	Volume file system analytics is not supported on SnapMirror destination volumes.
111411216	Enabling or disabling volume file system analytics is not supported on individual FlexGroup constituents.
111411217	Volume file system analytics is not supported on SnapLock volumes.
111411230	Volume file system analytics is not supported on volumes that contain NVMe namespaces.
111411241	Volume file system analytics is not supported for All SAN Array clusters.
111411252	Failed to enable file system analytics on volume "name" in SVM "svm.name" because there is insufficient available space. Ensure there is at least "space.available_percent`" available space in the volume, then try the operation again.
111411253	Failed to enable file system analytics on volume "name" in SVM "svm.name" because there is insufficient available space. Ensure there is at least "space.available_percent" available space in all constituents of the FlexGroup, then try the operation again."
111411257	Failed to enable file system analytics on volume "name" in SVM "svm.name" because there is insufficient available space. Increase the volume size to "size", then try the operation again."
124518405	Volume activity tracking is not supported on volumes that contain LUNs.
124518407	Volume activity tracking is not supported on FlexCache volumes.
124518408	Volume activity tracking is not supported on audit staging volumes.
124518409	Volume activity tracking is not supported on object store server volumes.
124518410	Volume activity tracking is not supported on SnapMirror destination volumes.
124518411	Enabling or disabling volume activity tracking is not supported on individual FlexGroup constituents.
124518412	Volume activity tracking is not supported on SnapLock volumes.

Error Code	Description
124518414	Volume activity tracking is not supported on volumes that contain NVMe namespaces.
124518422	Volume activity tracking is not supported on All SAN Array clusters.
144180203	Volume capacity rebalancing is not supported on FlexCache volumes.
144180204	Volume capacity rebalancing is not supported on object store volumes.
144180207	Volume capacity rebalancing is not supported on inactive MetroCluster configurations.
144182201	Volume capacity rebalancing using non-disruptive file move operations and granular data requires an effective cluster version of 9.11.1 or later.
144182203	The specified value for the "rebalancing.state" parameter is invalid for a PATCH operation. Valid values are "starting" and "stopping".
144182207	Modifying the volume capacity rebalancing configuration is not supported in the same operation that volume capacity rebalancing is being stopped.
144182211	The specified value for "-max-file-moves" is not valid.
144182212	The "-min-file-size" value specified must be larger than 0.
144182213	The "-min-threshold" value specified must be larger than 0.
144182214	The "-max-threshold" value specified must be larger than 0.
144182215	The maximum imbalance threshold value must be larger than the minimum imbalance threshold value.
144182216	Volume capacity rebalancing is running on the volume.
144182219	Volume capacity rebalancing for the volume has been started on some constituents.
144182221	The "-max-runtime" value specified must be 30 minutes or longer.
144182223	Volume capacity rebalancing is not running on the volume.
144182225	Internal error in the data component.
144182226	Failed to load the volume capacity rebalancing configuration for the volume.
144182233	The specified rebalancing start time must be set to the current time or a later time.

Error Code	Description
144182235	The "-min-file-size" value specified must be larger than or equal to 20MB.
196608022	Failed to start rekey on volume "name" in SVM "svm.name" because of the specified reason.
196608023	Failed to pause rekey on volume "name" because of the specified reason.
196608024	Failed to resume rekey on volume "name" because of the specified reason.
196608038	Failed to start conversion on volume "name" in SVM "svm.name" because of the specified reason.
196608039	Failed to pause conversion on volume "name" because of the specified reason.
196608041	Failed to resume conversion on volume "name" because of the specified reason.
196608207	Rekey operation is already in progress on volume "name".
196608208	Encryption conversion operation is already in progress on volume "name".
196608209	Operation is already paused.
203161623	Disabling anti-ransomware for volume "volume_name" in SVM "vserver_name" is in progress. Wait a few minutes, and try the operation again.
203161629	Anti-ransomware is not supported on a volume with Snapshot autodelete commitment set to "destroy".

Also see the table of common errors in the Response body overview section of this documentation.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

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. • enum: ["off", "on"] • Introduced in: 9.10 • x-nullable: true
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.

Name	Type	Description
unsupported_reason	unsupported_reason	

aggregates

Aggregate

Name	Type	Description
name	string	
uuid	string	

initialization

Name	Type	Description
state	string	State of the analytics file system scan.

scan_throttle_reason

Name	Type	Description
arguments	array[string]	Arguments present in the warning message encountered.
code	string	Warning code indicating why scanner throttling is reported.
message	string	A message that provides details for scan throttling.

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
files_scanned	integer	Number of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`.
initialization	initialization	
scan_progress	integer	Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is `initializing`.
scan_throttle_reason	scan_throttle_reason	
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 for all existing files and directories is currently running. If this value is "initialization_paused", this means that the initialization scan is currently paused. If this value is 'unknown', this means that there was an internal error when determining the file system analytics state for the volume. • enum: ["unknown", "initializing", "initialization_paused", "off", "on"] • Introduced in: 9.8 • x-nullable: true

Name	Type	Description
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.
total_files	integer	Total number of files in the volume that the file system analytics initialization scan will process. Only returned when the state is `initializing`.
unsupported_reason	unsupported_reason	

anti_ransomware_volume_attack_detection_parameters

Name	Type	Description
based_on_file_create_op_rate	boolean	Specifies whether attack detection is based on the file create operations rate. This parameter is valid only for NAS volumes.
based_on_file_delete_op_rate	boolean	Specifies whether attack detection is based on the file delete operations rate. This parameter is valid only for NAS volumes.
based_on_file_rename_op_rate	boolean	Specifies whether attack detection is based on the file rename operations rate. This parameter is valid only for NAS volumes.
based_on_high_entropy_data_rate	boolean	Specifies whether a high entropy data rate should be considered for attack detection.
based_on_never_seen_before_file_extension	boolean	Specifies whether file extensions never seen before should be considered for attack detection.

Name	Type	Description
file_create_op_rate_surge_notify_percent	integer	Specifies the percentage of surge in the file create rate up to which it is considered normal behavior.
file_delete_op_rate_surge_notify_percent	integer	Specifies the percentage of surge in the file delete rate up to which it is considered normal behavior.
file_rename_op_rate_surge_notify_percent	integer	Specifies the percent of surge in the file rename rate up to which it is considered normal behavior.
high_entropy_data_surge_notify_percent	integer	Specifies the percentage of surge in high entropy data up to which it is considered as normal behavior. For example, if the usual high entropy data rate in the volume is 5% and if this parameter is set to 100%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 10% at any time. Similarly, if this parameter is set to 400%, it will be considered as an unusual surge if the high entropy data rate of the volume exceeds 25%, and so on.
never_seen_before_file_extension_count_notify_threshold	integer	Specifies the number of files found with a never seen before file extension up to which it is considered normal behavior.
never_seen_before_file_extension_duration_in_hours	integer	Specifies the duration within which the specified number of files found with never seen before file extensions is considered normal behavior.
relaxing_popular_file_extensions	boolean	Specifies whether popular file extensions should be relaxed from being treated as a suspect for the attack. Some popular file extensions are .txt, .pdf, and so on.

anti_ransomware_attack_report

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

event_log

Name	Type	Description
is_enabled_on_new_file_extension_seen	boolean	Specifies whether to send an EMS when a new file extension is discovered.
is_enabled_on_snapshot_copy_creation	boolean	Specifies whether to send an EMS when a snapshot is created.

space

Name	Type	Description
snapshot_count	integer	Total number of Anti-ransomware backup snapshots.
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 snapshots.

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Peak rate of file creates per minute in the workload of the volume during surge.
file_delete_peak_rate_per_minute	integer	Peak rate of file deletes per minute in the workload of the volume during surge.

Name	Type	Description
file_rename_peak_rate_per_minute	integer	Peak rate of file renames per minute in the workload of the volume during surge.
high_entropy_data_write_peak_percent	integer	Peak percentage of high entropy data writes in the volume during surge.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Peak high entropy data write rate in the volume during surge, in KBs per minute.
time	string	Timestamp at which the first surge in the volume's workload is observed.

suspect_files

Name	Type	Description
count	integer	Total number of `suspect_files.format` files observed by the Anti-ransomware analytics engine on the volume.
entropy	string	Indicates the entropy level of this file type.
format	string	File formats observed by the Anti-ransomware analytics engine on the volume.

typical_usage

Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Typical peak rate of file creates per minute in the workload of the volume.
file_delete_peak_rate_per_minute	integer	Typical peak rate of file deletes per minute in the workload of the volume.

Name	Type	Description
file_rename_peak_rate_per_minute	integer	Typical peak rate of file renames per minute in the workload of the volume.
high_entropy_data_write_peak_percent	integer	Typical peak percentage of high entropy data writes in the volume.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Typical peak high entropy data write rate in the volume, in KBs per minute.

historical_statistics

Typical usage values of volume workload.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Typical peak rate of file creates per minute in the workload of the volume.
file_delete_peak_rate_per_minute	integer	Typical peak rate of file deletes per minute in the workload of the volume.
file_rename_peak_rate_per_minute	integer	Typical peak rate of file renames per minute in the workload of the volume.
high_entropy_data_write_peak_percent	integer	Typical peak percentage of high entropy data writes in the volume.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Typical peak high entropy data write rate in the volume, in KBs per minute.

newly_observed_file_extensions

Name	Type	Description
count	integer	Count of newly observed file extensions.
name	string	Name of the newly observed file extension.

surge_statistics

Usage values of the volume's workload during surge.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Peak rate of file creates per minute in the workload of the volume during surge.
file_delete_peak_rate_per_minute	integer	Peak rate of file deletes per minute in the workload of the volume during surge.
file_rename_peak_rate_per_minute	integer	Peak rate of file renames per minute in the workload of the volume during surge.
high_entropy_data_write_peak_percent	integer	Peak percentage of high entropy data writes in the volume during surge.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Peak high entropy data write rate in the volume during surge, in KBs per minute.
time	string	Timestamp at which the first surge in the volume's workload is observed.

surge_usage

Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.

Name	Type	Description
file_create_peak_rate_per_minute	integer	Peak rate of file creates per minute in the workload of the volume during surge.
file_delete_peak_rate_per_minute	integer	Peak rate of file deletes per minute in the workload of the volume during surge.
file_rename_peak_rate_per_minute	integer	Peak rate of file renames per minute in the workload of the volume during surge.

Name	Type	Description
high_entropy_data_write_peak_percent	integer	Peak percentage of high entropy data writes in the volume during surge.
high_entropy_data_write_peak_rate_kb_per_minute	integer	Peak high entropy data write rate in the volume during surge, in KBs per minute.
newly_observed_file_extensions	array[newly_observed_file_extensions]	New file extensions observed in the volume during surge.
time	string	Timestamp at which the first surge in the volume's workload is observed.

anti_ransomware_volume_workload

Name	Type	Description
file_extension_types_count	integer	Count of types of file extensions observed in the volume.
file_extensions_observed	array[string]	File extensions observed in the volume.
historical_statistics	historical_statistics	Typical usage values of volume workload.
newly_observed_file_extensions	array[newly_observed_file_extensions]	New file extensions observed in the volume during surge.
surge_statistics	surge_statistics	Usage values of the volume's workload during surge.
surge_usage	surge_usage	Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.
typical_usage	typical_usage	Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

anti_ransomware

Anti-ransomware related information of the volume.

Name	Type	Description
attack_detection_parameters	anti_ransomware_volume_attack_detection_parameters	
attack_probability	string	Probability of a ransomware attack. `none` No files are suspected of ransomware activity. `low` A number of files are suspected of ransomware activity. `moderate` A moderate number of files are suspected of ransomware activity. `high` 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.
event_log	event_log	
space	space	

Name	Type	Description
state	string	Anti-ransomware state. `disabled` Anti-ransomware monitoring is disabled on the volume. This is the default state in a POST operation. `disable_in_progress` Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. Valid in GET operation. `dry_run` Anti-ransomware monitoring is enabled in the evaluation mode. `enabled` Anti-ransomware monitoring is active on the volume. `paused` Anti-ransomware monitoring is paused on the volume. `enable_paused` Anti-ransomware monitoring is paused on the volume from its earlier enabled state. Valid in GET operation. `dry_run_paused` Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. Valid in GET operation. 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.
surge_as_normal	boolean	Indicates whether or not to set the surge values as historical values. This field is no longer supported. Use update_baseline_from_surge instead.
surge_usage	surge_usage	Usage values of the volume's workload during surge. This object is no longer supported use surge_statistics instead.
suspect_files	array[suspect_files]	
typical_usage	typical_usage	Typical usage values of volume workload. This object is no longer supported use historical_statistics instead.

Name	Type	Description
update_baseline_from_surge	boolean	Sets the observed surge value as the new baseline on a volume.
workload	anti_ransomware_volume_workload	

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.

asynchronous_directory_delete

Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.

Name	Type	Description
enabled	boolean	Specifies whether asynchronous directory delete from the client is enabled on the volume.
trash_bin	string	Name of the trash bin directory. If no "trash_bin" property is specified when enabling, the default trash bin name, "._ontaptrashbin", is used.

autosize

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

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

snapshot_reference

Name	Type	Description
name	string	
uuid	string	

parent_svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
name	string	The name of the SVM. This field cannot be specified in a PATCH method.
uuid	string	The unique identifier of the SVM. This field cannot be specified in a PATCH method.

parent_volume

Name	Type	Description
name	string	The name of the volume. This field cannot be specified in a PATCH method.

Name	Type	Description
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6 • x-nullable: true

clone

Name	Type	Description
has_flexclone	boolean	Specifies whether this volume is a parent of any FlexClone volume.
inherited_physical_used	integer	Inherited physical used from the clone's base snapshot.
inherited_savings	integer	Inherited savings from the clone's base snapshot.
split_complete_percent	integer	Percentage of FlexClone volume 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 a split is executed on any FlexClone volume, that is when the FlexClone volume is split from its parent FlexVol volume. Setting this field initiates a split of a FlexClone volume from a FlexVol volume. 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 volume constituent.
uuid	string	Unique identifier for the aggregate.

destination_aggregate

Aggregate

Name	Type	Description
name	string	
uuid	string	

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup volume constituents is 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 autodelete is currently enabled on this volume.
reserve_percent	integer	The space that has been set aside as a reserve for snapshot usage, in percent.
used	integer	The total space used by snapshots 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.
large_size_enabled	boolean	Specifies whether the support for large volumes and large files is enabled on the volume.
local_tier_footprint	integer	Space used by the local tier for this volume in the aggregate, in bytes.
logical_space	logical_space	
max_size	string	Indicates the maximum size supported for each constituent of the FlexGroup volume.
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.
total_metadata	integer	Space used by the total metadata in the volume, in bytes.
total_metadata_footprint	integer	Space used by the volume metadata footprint in the aggregate, in bytes.
used	integer	The virtual space used (includes volume reserves) before storage efficiency, in bytes.

Name	Type	Description
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	
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 volume constituents is not supported.
name	string	FlexGroup volume constituent name.
space	space	

idcs_scanner

Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner.

Name	Type	Description
enabled	boolean	Specifies the administrative state of the inactive data compression scanner. Disabling inactive data compression is not allowed on Capacity optimized Flash with QAT supported platforms.
inactive_days	integer	Data blocks older than, or equal to, 'inactive_days' are picked up by the inactive data compression scanner. Valid for PATCH only. Only applicable when 'operation_state' set to 'active'.

Name	Type	Description
mode	string	Specifies the mode of inactive data compression scanner. Valid for PATCH and GET.
operation_state	string	Specifies the operational state of the inactive data compression scanner. VALID for PATCH and GET. Valid options for PATCH are "idle" and "active".
status	string	Status of last inactive data compression scan on the volume.
threshold_inactive_time	string	Time interval after which inactive data compression is automatically triggered. The value is in days and is represented in the ISO-8601 format "P<num>D", for example "P3D" represents a duration of 3 days. This field is not supported on QAT supported platforms.</num>

policy

Name	Type	Description
name	string	Specifies the name of the efficiency policy. The "inline-only" and "none" policies are not supported on Capacity optimized Flash with QAT supported platform.

scanner

Name	Type	Description
compression	boolean	Start compression if scanning old data. Valid for PATCH and GET. This option is not supported for FSX/CVO platforms.
dedupe	boolean	Start deduplication if scanning old data. Valid for PATCH and GET.

Name	Type	Description
scan_old_data	boolean	Indicates whether or not to scan old data. Valid for PATCH and GET.
state	string	State of the volume efficiency scanner. Valid for PATCH and GET. Valid options for PATCH are "idle" and "active".

space_savings

Name	Type	Description
compression	integer	Total disk space that is saved by compressing blocks on the referenced file system, in bytes.
compression_percent	integer	Percentage of total disk space that is saved by compressing blocks on the referenced file system.
dedupe	integer	Total disk space that is saved by deduplication and file cloning, in bytes.
dedupe_percent	integer	Percentage of total disk space that is saved by deduplication and file cloning.
dedupe_sharing	integer	Total disk space that is shared due to deduplication and file cloning.
total	integer	Total disk space saved in the volume due to deduplication, compression and file cloning, in bytes.
total_percent	integer	Percentage of total disk space saved in the volume due to deduplication, compression and file cloning.

efficiency

Name	Type	Description
application_io_size	string	Block size to use by compression. 8k and auto are only allowed for POST. Only auto is supported on QAT supported platforms.
auto_state	string	Automatic deduplication schedule volume state. auto ‐ Volumes with auto_state set to auto start post-process deduplication automatically. deprioritized ‐ Volumes with auto_state set to deprioritized do not start post-process deduplication automatically.
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 FlexGroup volumes, where some of the constituent volumes are compaction enabled and some are disabled.

Name	Type	Description
compression	string	The system can be enabled/disabled compression. Disabling compression is not allowed on Capacity optimized Flash with QAT supported platforms. inline ‐ Data will be compressed first and written to the volume. background ‐ Data will be written to the volume and compressed later. both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run. none ‐ None mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are compression enabled and some are disabled. that On volumes with container compression enabled, background compression refers to inactive data compression scan enabled on the volume.
compression_type	string	Compression type to use by compression. Valid for PATCH and GET.

Name	Type	Description
cross_volume_dedupe	string	The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. Disabling cross volume dedupe is not allowed on Capacity optimized Flash with QAT supported platforms. inline ‐ Data will be cross volume deduped first and written to the volume. background ‐ Data will be written to the volume and cross volume deduped later. both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroup volumes, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.
dedupe	string	The system can be enabled/disabled dedupe. Disabling dedupe is not allowed on Capacity optimized Flash with QAT supported platforms. 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 FlexGroup volumes, where some of the constituent volumes are dedupe enabled and some are disabled.
has_savings	boolean	When true, indicates that the volume contains shared(deduplication, file clones) or compressed data.

Name	Type	Description
idcs_scanner	idcs_scanner	Inactive data compression scan looks and picks up blocks that have not been read for a certain amount of time(threshold_inactive_days). These blocks are then compressed in 32K chunks. All attributes are valid for GET only, except for 'operation_state' which is valid for PATCH and GET, and is used to start/stop the scanner.
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.
logging_enabled	boolean	When true, indicates that space savings for any newly-written data are being logged.
op_state	string	Sis status of the volume.
policy	policy	
progress	string	Sis progress of the volume.
scanner	scanner	
schedule	string	Schedule associated with volume.
space_savings	space_savings	

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 and Cloud Volumes 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 and Cloud Volumes ONTAP ‐ enabled ‐ All supported storage efficiency features for the volume are enabled. custom ‐ Read-only field currently only supported for the FSx for ONTAP and Cloud Volumes ONTAP, user-defined storage efficiency features are enabled. For other platforms ‐ enabled ‐ At least one storage efficiency feature for the volume is enabled. • enum: ["disabled", "enabled", "mixed", "custom"] • Introduced in: 9.9 • x-nullable: true
storage_efficiency_mode	string	Storage efficiency mode used by volume. This parameter is supported only on AFF platforms. There is no difference between default and efficient modes on QAT supported platforms and auto adaptive compression is set irrespective of the modes.
type	string	Sis Type of the volume.
volume_path	string	Absolute volume path of the volume.

status

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

encryption

Name	Type	Description
action	string	This field can be used to pause an ongoing rekey or conversion operation or resume a paused rekey or conversion operation. Valid in PATCH. 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.
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_create_time	string	Encryption key creation time of the volume.

Name	Type	Description
key_id	string	The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key.
rekey	boolean	If set to 'true', re-encrypts the volume with a new key. Valid in PATCH.
state	string	Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one.
status	status	
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.

Name	Type	Description
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.

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

throughput

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

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

metric

Performance numbers, such as IOPS, latency and throughput.

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

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

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 volume 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
id	integer	
name	string	

junction_parent

Name	Type	Description
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 dight), first digit is assumed to be zero.

policy

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

Name	Type	Description
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

Name	Type	Description
min_throughput_iops	integer	Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_mbps	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
name	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
uuid	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

qos

QoS information

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

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume 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 FlexVol volumes or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVol volumes or FlexGroup volumes when the quota state is "off".
state	string	Quota state of the volume

last_error

Error information for the last failed file move on the constituent.

Name	Type	Description
code	integer	Error code of the last file move error on the constituent.
destination	integer	DSID of the destination constituent of the last file move error on the constituent.
file_id	integer	File ID of the last file move error on the constituent.
time	string	Time of the last file move error on the constituent.

movement

Properties on this constituent related to file movement.

Name	Type	Description
file_moves_started	integer	Number of file moves started on this constituent.
last_error	last_error	Error information for the last failed file move on the constituent.

Name	Type	Description
most_recent_start_time	string	Start time of the most recent file move on the constituent.

blocks_skipped

Number of blocks skipped by the scanner on this constituent due to various reasons.

Name	Type	Description
efficiency_blocks	integer	Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.
efficiency_percent	integer	Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.
fast_truncate	integer	Number of blocks skipped by the scanner on this constituent because fast truncate is currently running on files.
footprint_invalid	integer	Number of blocks skipped by the scanner on this constituent because of files with invalid space footprints.
in_snapshot	integer	Number of blocks skipped by the scanner on this constituent because of files in snapshots.
incompatible	integer	Number of blocks skipped by the scanner on this constituent because of incompatible files.
metadata	integer	Number of blocks skipped by the scanner on this constituent because of metadata files.
on_demand_destination	integer	Number of blocks skipped by the scanner on this constituent because of on demand destination files.

Name	Type	Description
other	integer	Number of blocks skipped by the scanner on this constituent for all other reasons.
remote_cache	integer	Number of blocks skipped by the scanner on this constituent because of remote caches.
too_large	integer	Number of blocks skipped by the scanner on this constituent because of files that are larger than rebalancing.max_file_size.
too_small	integer	Number of blocks skipped by the scanner on this constituent because of files that are smaller than rebalancing.min_file_size.
write_fenced	integer	Number of blocks skipped by the scanner on this constituent because of files fenced for write operations.

files_skipped

Number of files skipped by the scanner on this constituent due to various reasons.

Name	Type	Description
efficiency_blocks	integer	Number of files skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.
efficiency_percent	integer	Number of files skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.
fast_truncate	integer	Number of files skipped by the scanner on this constituent because fast truncate is currently running on the file.
footprint_invalid	integer	Number of files skipped by the scanner on this constituent because their space footprints are invalid.

Name	Type	Description
in_snapshot	integer	Number of files skipped by the scanner on this constituent because they are trapped in snapshots.
incompatible	integer	Number of files skipped by the scanner on this constituent because they are incompatible.
metadata	integer	Number of files skipped by the scanner on this constituent because they metadata files.
on_demand_destination	integer	Number of files skipped by the scanner on this constituent because they are on demand destinations.
other	integer	Number of files skipped by the scanner on this constituent for all other reasons.
remote_cache	integer	Number of files skipped by the scanner on this constituent because they are remote caches.
too_large	integer	Number of files skipped by the scanner on this constituent because they are larger than rebalancing.max_file_size.
too_small	integer	Number of files skipped by the scanner on this constituent because they are smaller than rebalancing.min_file_size.
write_fenced	integer	Number of files skipped by the scanner on this constituent because they are fenced for write operations.

scanner

Properties related to determining which files to move and where to move them to.

Name	Type	Description
blocks_scanned	integer	Number of blocks scanned on this constituent.
blocks_skipped	blocks_skipped	Number of blocks skipped by the scanner on this constituent due to various reasons.
files_scanned	integer	Number of files scanned on this constituent.
files_skipped	files_skipped	Number of files skipped by the scanner on this constituent due to various reasons.

engine

Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions.

Name	Type	Description
movement	movement	Properties on this constituent related to file movement.
scanner	scanner	Properties related to determining which files to move and where to move them to.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

rebalancing

Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

Name	Type	Description
data_moved	integer	The amount of data that has been moved in or out of a constituent. A positive value represents data moving into the constituent while a negative value is data moving out of the constituent.
engine	engine	Each constituent has one rebalancing engine that co-ordinates scanning constituents for free space and files to move, as well as moving files between constituents. The engine properties must be explicitly requested, are meant for helping diagnose why the volume rebalancer is making decisions.
exclude_snapshots	boolean	Specifies whether or not to exclude files that are stuck in snapshots during rebalancing operation. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "exclude_snapshots" value. Once the operation is started, any changes to the "exclude_snapshots" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "exclude_snapshots" value.
imbalance_percent	integer	Represents the percentage the volume is out of balance.
imbalance_size	integer	Represents how much the volume is out of balance, in bytes.

Name	Type	Description
max_constituent_imbalance_percent	integer	Absolute percentage of the constituent that is most out of balance. This value will update every 30 seconds when rebalancing is not active and every 10 seconds when rebalancing is active.
max_file_moves	integer	Specifies the maximum number of concurrent file moves in a volume capacity rebalancing operation on a constituent of the FlexGroup volume. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "max_file_moves" value. Once the operation is started, any changes to the "max_file_moves" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "max_file_moves" value.
max_runtime	string	This optional field specifies the maximum time a capacity rebalancing operation runs for. Once the maximum runtime has passed, the capacity rebalancing operation stops. If it is not set, the default value is 6 hours. This value cannot be updated while a capacity rebalancing operation is running. The maximum runtime 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 "P3D" represents a duration of 3 days. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively.</num></num></num></num></num>

Name	Type	Description
max_threshold	integer	Specifies the maximum imbalance percentage for FlexGroup volume constituents. When a constituent's imbalance percentage is larger than this value, files are moved from the constituent. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "max_threshold" value. Once the operation is started, any changes to the "max_threshold" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "max_threshold" value.
min_file_size	integer	Specifies the minimum file size to consider for a volume capacity rebalancing operation. When a new capacity rebalancing operation is started on a FlexGroup volume, it uses the current "min_file_size" value. Once the operation is started, any changes to the "min_file_size" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "min_file_size" value. The value must be a multiple of 4KB. If it is not set, the default value is 100MB. Setting "min-file-size" to less than the default value leads to more files being moved. Moved files use granular data, which may impact read/write I/O performance.

Name	Type	Description
min_threshold	integer	Specifies the minimum imbalance percentage for FlexGroup volume constituents. When a constituent's imbalance percentage is smaller than this value, files are not moved from the constituent. When a new capacity rebalancing operation is started on a FlexGroup volume, it will use the current "min_threshold" value. Once the operation is started, any changes to the "min_threshold" value do not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new "min_threshold" value.
notices	array[error]	Capacity rebalancing notice messages.
runtime	string	Duration the capacity rebalancing operation has been running.
start_time	string	Time when the current capacity rebalancing operation started, or when a future scheduled rebalancing operation begins.

Name	Type	Description
state	string	State of the volume capacity rebalancing operation. PATCH the state to "starting" to trigger the capacity rebalance operation, and include start_time to schedule rebalancing. PATCH the state to "stopping" to stop the capacity rebalance operation, or cancel a scheduled rebalancing operation. PATCH without the state with a valid start_time to modify the start_time of an existing scheduled rebalance operation. While a FlexGroup volume is rebalancing, every constituent will have a rebalancing engine that can either be scanning the filesystem for space usage and files to move, actively moving files or temporarily doing neither. If one or more constituents has a state of "rebalancing_source" or "rebalancing_dest", then files are being moved to rebalance the FlexGroup volume. If no files are being moved, more information about what the rebalancing engine is doing for each constituent is available using the "rebalancing.engine" property. The following values apply to FlexGroup volumes. not_running ‐ capacity rebalancing is not running on the volume. starting ‐ used in a PATCH operation to start a capacity rebalancing operation. rebalancing ‐ capacity rebalancing is running on the volume. paused ‐ volume capacity rebalancing is paused on the volume. stopping ‐ used in a PATCH operation to stop a capacity rebalancing operation. unknown ‐ the system was unable to determine the rebalancing state for the volume. The following values apply to FlexGroup volume constituents.

Name	Type	Description
stop_time	string	Time when the capacity rebalancing operation stopped.
target_used	integer	Represents the ideal used size of each constituent. Calculated by dividing the total FlexGroup volume used size by the number of constituents.
used_for_imbalance	integer	Represents the used size of each constituent, as determined by the rebalancing engine. Calculated by subtracting the size used by snapshots, the size of files pending deletion and the size of filesystem metadata from the volume used size.

retention

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

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

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

snaplock

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

Name	Type	Description
autocommit_period	string	Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".</num></num></num></num></num>
compliance_clock_time	string	This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
expiry_time	string	Expiry time of the volume. For volumes with an infinite, indefinite or unspecified SnapLock expiry time, "9999-12-31T23:59:59" is used to denote the time.
is_audit_log	boolean	Indicates if this volume has been configured as SnapLock audit log volume for the SVM .

Name	Type	Description
litigation_count	integer	Litigation count indicates the number of active legal-holds on the volume.
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	
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. • readOnly: 1 • Introduced in: 9.9 • x-nullable: true

snapmirror

Specifies attributes for SnapMirror protection.

Name	Type	Description
destinations	destinations	

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

snapshot_policy

This is a reference to the snapshot policy.

Name	Type	Description
name	string	
uuid	string	

logical_space

Name	Type	Description
available	integer	The amount of space available in this volume with storage efficiency space considered used, in bytes.
enforcement	boolean	Specifies whether space accounting for operations on the volume is done along with storage efficiency.
reporting	boolean	Specifies whether space reporting on the volume is done along with storage efficiency.
used	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.
used_by_snapshots	integer	Size that is logically used across all snapshots in the volume, in bytes.

Name	Type	Description
used_percent	integer	SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), as a percentage.

autodelete

Name	Type	Description
commitment	string	By default, snapshot autodelete does not delete snapshots locked by Snapmirror, clones of a volume, a LUN, an NVMe namespace, or a file. Deletion of snapshots locked by these applications is specified using this option. The default value is try.
defer_delete	string	Allows the user to inform snapshot autodelete to defer the deletion of a specified snapshot until the end. The default value is user_created.
delete_order	string	Specifies the order in which snapshot autodelete occurs. Ordering is done using the date and time the snapshot is created. The default value is oldest_first.
enabled	boolean	Specifies whether snapshot autodelete is currently enabled on this volume.
prefix	string	Specifies the prefix of the snapshot which if matched, is deleted last. Used with autodelete_defer_delete when used with a prefix value.
target_free_space	integer	Snapshots are deleted, one at a time, until the used volume space reaches the value specified. The default is 20% free space or 80% utilized.

Name	Type	Description
trigger	string	Specifies when the system should trigger an autodelete of snapshots. When set to *volume*, autodelete is triggered based on volume fullness. When set to *snap_reserve*, autodelete is triggered based on snapshot reserve fullness. The default value is *volume*.

snapshot

Name	Type	Description
autodelete	autodelete	
autodelete_enabled	boolean	Specifies whether snapshot autodelete is currently enabled on this volume. This field will no longer be supported in a future release. Use autodelete.enabled instead.
autodelete_trigger	string	Specifies when the system should trigger an autodelete of snapshots. When set to *volume*, autodelete is triggered based on volume fullness. When set to *snap_reserve*, autodelete is triggered based on snapshot reserve fullness. The default value is *volume*. This field will no longer be supported in a future release. Use autodelete.trigger instead.
reserve_available	integer	Size available for snapshots within the snapshot reserve, in bytes.
reserve_percent	integer	The space that has been set aside as a reserve for snapshot usage, in percent.
reserve_size	integer	Size in the volume that has been set aside as a reserve for snapshot usage, in bytes.

Name	Type	Description
space_used_percent	integer	Percentage of snapshot reserve size that has been used.
used	integer	The total space used by snapshots in the volume, in bytes.

space

Name	Type	Description
afs_total	integer	Total size of AFS, excluding snap-reserve, in bytes.
auto_adaptive_compression_footprint_data_reduction	integer	Savings achieved due to Auto Adaptive Compression, 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.
capacity_tier_footprint_data_reduction	integer	Savings achieved in the space used by the capacity tier for this volume in the FabricPool aggregate, in bytes.
compaction_footprint_data_reduction	integer	Savings achieved due to Data Compaction, in bytes.

Name	Type	Description
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.
effective_total_footprint	integer	Volume footprint after efficiency savings, in bytes. effective total footprint represents total footprint after deducting auto adaptive compression and compaction savings. effective-footprint includes aggregate metadata used by volume.
expected_available	integer	Size that should be available for the volume, irrespective of available size in the aggregate, in bytes.
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.

Name	Type	Description
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.
is_used_stale	boolean	Specifies if the virtual space used is stale.
large_size_enabled	boolean	Indicates if the support for large FlexVol volumes and large files is enabled on this volume. When configured to true, FlexVol volume size can reach up to 300TB and single file size can reach 128TB.
local_tier_footprint	integer	Space used by the local tier for this volume in the aggregate, in bytes.
logical_space	logical_space	
max_size	string	Indicates the maximum size supported for the FlexVol volume or for each constituent of the FlexGroup volume.
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 is effective total footprint which is equivalent to total footprint after deducting auto adaptive compression and compaction savings. Physical-used also includes aggregate metadata used by volume.
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 snapshots 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 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.
total_metadata	integer	Space used by the total metadata in the volume, in bytes.

Name	Type	Description
total_metadata_footprint	integer	Space used by the volume metadata footprint 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.
volume_guarantee_footprint	integer	Space reserved for future writes in the volume, in bytes.

access

Raw count and latency data for access operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

audit

Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

dir

Raw count and latency data for directory-create operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

file

Raw count and latency data for file-create operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

other

Raw count and latency data for create operations on objects other than files, directories and symlinks.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

symlink

Raw count and latency data for symlink-create operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

create

Raw count and latency data for create operations.

getattr

Raw count and latency data for getattr operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

link

Raw count and latency data for link operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

lock

Raw count and latency data for lock operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

lookup

Raw count and latency data for lookup operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

open

Raw count and latency data for open operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

read

Raw count and latency data for read operations, including histograms categorizing operations by size and latency.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • readOnly: 1 • x-ntap-advanced: true • Introduced in: 9.11 • x-nullable: true
volume_protocol_latency_histogram_counts	array[integer]	
volume_protocol_latency_histogram_labels	array[string]	Labels for the latency histogram, ranging from <2us to >20s.
volume_protocol_size_histogram_counts	array[integer]	
volume_protocol_size_histogram_labels	array[string]	Labels for the size histogram, ranging from <4KB to >1024KB.

readdir

Raw count and latency data for readdir operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

readlink

Raw count and latency data for readlink operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

rename

Raw count and latency data for rename operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

setattr

Raw count and latency data for setattr operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

unlink

Raw count and latency data for unlink operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

watch

Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.

Name	Type	Description
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • Introduced in: 9.11 • x-nullable: true

write

Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

Name	Type	Description
count	integer	Number of operations of the given type performed on this volume.
total_time	integer	The raw data component latency in microseconds measured within ONTAP for all operations of the given type. • example: 200 • readOnly: 1 • x-ntap-advanced: true • Introduced in: 9.11 • x-nullable: true
volume_protocol_latency_histogram_counts	array[integer]	
volume_protocol_latency_histogram_labels	array[string]	Labels for the latency histogram, ranging from <2us to >20s.
volume_protocol_size_histogram_counts	array[integer]	
volume_protocol_size_histogram_labels	array[string]	Labels for the size histogram, ranging from <4KB to >1024KB.

cifs_ops_raw

Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

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.

nfs_ops_raw

Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.

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
cifs_ops_raw	cifs_ops_raw	Raw data component performance values for CIFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.
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.

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.
nfs_ops_raw	nfs_ops_raw	Raw data component performance values for NFS operations on this volume, including number of operations and raw latency, in microseconds for each operation. The values are cumulative and increase while the volume is online.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_ delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

Name	Type	Description
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
name	string	The name of the SVM. This field cannot be specified in a PATCH method.
uuid	string	The unique identifier of the SVM. This field cannot be specified in a PATCH method.

tiering

Name	Type	Description
min_cooling_days	integer	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.
object_tags	array[string]	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".

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

volume

Name	Type	Description
_tags	array[string]	Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings.
access_time_enabled	boolean	Indicates whether or not access time updates are enabled on the volume.
activity_tracking	activity_tracking	

Name	Type	Description
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
aggressive_readahead_mode	string	Specifies the `aggressive_readahead_mode` enabled on the volume. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. When the option is set to *sequential_read*, the system aggressively prefetches the file completely, or to a certain length based on the file size limit, and continues as the read makes progress. If the option is set to *cross_file_sequential_read*, then the system aggressively prefetches multiple files completely, or to a certain length, and continues as the read makes progress.
analytics	analytics	
anti_ransomware	anti_ransomware	Anti-ransomware related information of the volume.
application	application	
asynchronous_directory_delete	asynchronous_directory_delete	Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.
autosize	autosize	
clone	clone	

Name	Type	Description
cloud_retrieval_policy	string	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.
cloud_write_enabled	boolean	Indicates whether or not cloud writes are enabled on the volume. NFS writes to this volume are sent to the cloud directly instead of the local performance tier. This feature is only available on volumes in FabricPools.
comment	string	A comment for the volume. Valid in POST or PATCH.
consistency_group	consistency_group	Consistency group the volume is part of.

Name	Type	Description
constituents	array[constituents]	FlexGroup volume constituents. FlexGroup volume constituents can be retrieved more efficiently by specifying "is_constituent=true" or "is_constituent=true&flexgroup.uuid=<flexgroup.uuid>" as query parameters.</flexgroup.uuid>
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. If a volume is being created on multiple aggregates and the "constituents_per_aggregate" field is not specified, the default value of the "constituents_per_aggregate" field is 4. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list, unless 'optimize_aggregates' is specified as 'true'. If the "aggregates.name" or "aggregates.uuid" is specified in a PATCH request to expand an existing FlexGroup volume, the default value of the "constituents_per_aggregate" field is 1. The volume expand operation is only supported on FlexGroup volumes.
convert_unicode	boolean	Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients.

Name	Type	Description
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	
granular_data	boolean	State of granular data on the volume. This setting is `true` by default when creating an AWS S3 FlexGroup volume via a POST request and `false` by default for creating any other type of FlexGroup volume via a POST request. On FlexVol volumes, the setting is always `false`, as only FlexGroup volumes support this feature. Once enabled, this setting can only be disabled by restoring a snapshot. Earlier versions of ONTAP (pre 9.12) are not compatible with this feature. Therefore, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the setting was enabled. • Introduced in: 9.12 • x-nullable: true

Name	Type	Description
granular_data_mode	string	Mode of granular data on the volume. This setting defaults to `basic` when the `granular_data` parameter is set to `true`, but can be specified at the time of creation via a POST request. Earlier versions of ONTAP (pre 9.12) are not compatible with the `basic` setting. Therefore, when set to `basic`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `basic` mode was enabled. Earlier versions of ONTAP (pre 9.16) are not compatible with the `advanced` setting. Therefore, when set to `advanced`, reverting to an earlier version of ONTAP is not possible unless this volume is deleted or restored to a snapshot that was taken before the `advanced` mode was enabled. • enum: ["disabled", "basic", "advanced"] • Introduced in: 9.16 • x-nullable: true
guarantee	guarantee	
is_object_store	boolean	Specifies whether the volume is provisioned for an object store server.
is_svm_root	boolean	Specifies whether the volume is a root volume of the SVM it belongs to.

Name	Type	Description
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.
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 volume constituents are not supported.
msid	integer	The volume's Master Set ID.
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 FlexGroup volumes, 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.
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol volume or a FlexGroup volume.

Name	Type	Description
rebalancing	rebalancing	Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.
scheduled_snapshot_naming_scheme	string	Naming Scheme for automatic snapshots: • create_time - Automatic snapshots 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 snapshots in the volume.
snapshot_directory_access_enabled	boolean	This field, if true, enables the visible ".snapshot" directory from the client. The ".snapshot" directory will be available in every directory on the volume.
snapshot_locking_enabled	boolean	Specifies whether or not snapshot copy locking is enabled on the volume.

Name	Type	Description
snapshot_policy	snapshot_policy	This is a reference to the snapshot policy.
space	space	
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.
status	array[string]	Describes the current status of a volume.

Name	Type	Description
style	string	The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates or license. 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 volume. When the UDO License is installed, and no aggregates are specified, the system automatically provisions a FlexGroup volume on system selected aggregates. 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 ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroup volumes flexgroup_constituent ‐ FlexGroup volume constituents.
tiering	tiering	
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6 • x-nullable: true

Name	Type	Description
validate_only	boolean	Validate the volume move or volume conversion operations and their parameters, without actually performing the operation.

job_link

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

returned_error

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

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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