



Storage

ONTAP 9.13.1 REST API reference

NetApp
April 02, 2024

Table of Contents

- Storage 1
 - Storage overview 1
 - Retrieve or create a collection of storage aggregates 1
 - Retrieve a collection of cloud stores used by an aggregate 116
 - Attach an object store to an aggregate or add a second object store as a mirror 123
 - Remove a cloud target from an aggregate 132
 - Retrieve the cloud store for an aggregate 135
 - Update a cloud store 140
 - Manage storage aggregate plexes 147
 - Manage storage aggregates 167
 - Storage aggregate metrics 273
 - Manage bridges 287
 - Manage disks 326
 - Manage storage file clones 386
 - Manage storage file moves 458
 - Manage FlexCache volumes 528
 - Manage storage FlexCache origin volumes 589
 - Manage storage pools 618
 - Manage specific storage pools 650
 - Manage storage ports 687
 - Manage QoS policies 714
 - Manage QoS workloads 751
 - Manage storage qtrees 763
 - Storage quota 833
 - Manage storage quota reports 835
 - Manage storage quota policy rules 877
 - Manage shelves 942
 - Create and retrieve Snapshot copy policies 1013
 - Manage Snapshot copy policies and schedules 1033
 - Manage storage switches 1074
 - Manage storage tape devices 1103
 - Manage volume efficiency policies 1137

Storage

Storage overview

Overview

The ONTAP storage APIs can be used to manage physical and logical storage. This includes management of aggregates, volumes, LUNs, qtrees, snapshots, quotas, and storage efficiency.

Retrieve or create a collection of storage aggregates

Storage aggregates endpoint overview

Retrieving storage aggregate information

The Storage Aggregate GET API retrieves all data aggregates in the cluster. System owned root aggregates are not included in the output. This API also supports specific queries, in addition to queries on aggregate body properties, which affect the output of the API. The parameters for these queries are "recommend" and "show_spare". Using the "recommend" query returns the list of aggregates that are recommended for creation in the cluster. The "show_spare" query returns a response outside of the records body, which includes the groups of usable spares in the cluster. The "show_spare" query along with "flash_pool_eligible" restricts groups of usable spares for whole disk based flash pool creation. For storage pool creation, the appropriate groups of eligible spares will be returned by a GET on storage/pools?show_spare=true. The usable count for each class of spares does not include reserved spare capacity recommended by ONTAP best practices.

The collection GET returns the aggregate identifiers, UUID and name, and the node on which the aggregate resides. The instance GET, by default, returns all of the properties defined in the aggregates object, except advanced properties. The properties "space.footprint", "is_spare_low", "space.block_storage.inactive_user_data", "space.block_storage.inactive_user_data_percent", "inode_attributes.version", "inode_attributes.files_private_used", "inode_attributes.file_private_capacity" and "inode_attributes.file_public_capacity" are considered advanced properties and only returned when requested using the "fields" query parameter. Performance "metric" and "statistics" for aggregates are also only returned when requested. The "statistics" property accounts for the cumulative raw values collected by ONTAP for an aggregate, while the "metric" property displays the incremental average for latency and incremental changes in IOPs and throughput over the last 15 seconds. Any external application can use the raw statistics to derive its own incremental performance metrics. The collection GET does not support queries in conjunction with the "order_by" parameter.

Creating storage aggregates

When the POST command is issued with no properties, the system evaluates the cluster attached storage, determines the optimal aggregate layout and configures the aggregates. This layout is completely controlled by the system. To view the recommended optimal layout rather than creating it, use the GET endpoint, setting the "recommend" query to 'true'. This response will also include any warnings related to any inefficiencies in the storage configuration. Recommended aggregate creation is not supported on ONTAP Cloud and MetroCluster with Fibre Channel (FC). Alternatively, POST can be used with specific properties to create an aggregate as requested. At a minimum, the aggregate name, disk count, and the node where it should reside are required if any properties are provided.

When using POST with input properties, three properties are required. These are:

- name

- node.name or node.uuid
- block_storage.primary.disk_count

Remaining properties are optional

The following properties can be specified in POST:

- name - Name of the aggregate.
- node.name and node.uuid - Node on which the aggregate will be created.
- block_storage.primary.disk_count - Number of disks to be used to create the aggregate.
- block_storage.mirror.enabled - Specifies whether or not the aggregate should be created using SyncMirror.
- block_storage.primary.checksum_style - Checksum style of the disks to be use for the aggregate.
- block_storage.primary.disk_class - Class of disks to be use to for the aggregate.
- block_storage.primary.raid_size - Desired RAID size of the aggregate.
- block_storage.primary.raid_type - Desired RAID type of the aggregate.
- snaplock_type - SnapLock type to use on the aggregate.
- data_encryption.software_encryption_enabled - Enable or disable NAE (NetApp Aggregate Encryption) on the aggregate.
- simulate - Simulate the creation of the aggregate with specified input parameters. If the simulate field is specified, the response includes information on the proposed aggregate disk layout, any associated warnings, along with the proposed final size of the aggregate.

Examples

Retrieving a list of aggregates from the cluster

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

```

# The API:
/api/storage/aggregates

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

# The response:
{
  "records": [
    {
      "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
      "name": "test1",
      "node": {
        "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
        "name": "node-1",
      },
    },
    {
      "uuid": "4a7e4139-ca7a-420b-9a11-3f040d2189fd",
      "name": "test4",
      "node": {
        "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
        "name": "node-2",
      },
    }
  ],
  "num_records": 2,
}

```

Retrieving a list of aggregates recommended for creation from the cluster

The following example shows the response with a list of recommended data aggregates in the cluster.



Each aggregate UUID provided in this response is not guaranteed to be the same UUID for the aggregate if it is created.

```

# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/aggregates?recommend=true&fields=*" -H "accept:
application/json"

```

```

# The response:
{
  "records": [
    {
      "uuid": "795bf7c2-fa4b-11e8-ba65-005056bbe5c1",
      "name": "node_2_SSD_1",
      "node": {
        "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
        "name": "node-2",
      },
      "space": {
        "block_storage": {
          "size": 1116180480
        }
      },
      "block_storage": {
        "primary": {
          "disk_count": 22,
          "disk_class": "solid_state",
          "raid_type": "raid_dp",
          "disk_type": "ssd",
          "raid_size": 24,
          "simulated_raid_groups": [
            {
              "name": "node_2_SSD_1/plex0/rg0",
              "raid_type": "raid_dp",
              "parity_disk_count": 2,
              "data_disk_count": 9,
              "usable_size": 12309487,
              "is_partition": true
            },
            {
              "name": "node_2_SSD_1/plex1/rg0",
              "raid_type": "raid_dp",
              "parity_disk_count": 2,
              "data_disk_count": 9,
              "usable_size": 12309487,
              "is_partition": false
            }
          ]
        },
        "hybrid_cache": {
          "enabled": true,
          "storage_pools": [
            {
              "allocation_units_count": 1,

```

```

        "storage_pool": {
            "name": "sp1",
            "uuid": "1511d084-7290-11ec-ae5b-005056bb2afa"
        }
    },
    {
        "allocation_units_count": 1,
        "storage_pool": {
            "name": "sp2",
            "uuid": "342d234f-7291-11ec-ae5b-005056bb2afa"
        }
    }
]
},
"mirror": {
    "enabled": true
}
},
{
    "uuid": "795c0a15-fa4b-11e8-ba65-005056bbe5c1",
    "name": "node_1_SSD_1",
    "node": {
        "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
        "name": "node-1",
    },
    "space": {
        "block_storage": {
            "size": 352477184
        }
    },
    "block_storage": {
        "primary": {
            "disk_count": 22,
            "disk_class": "solid_state",
            "raid_type": "raid_dp",
            "disk_type": "ssd",
            "raid_size": 24,
            "simulated_raid_groups": [
                {
                    "name": "node_1_SSD_1/plex0/rg0",
                    "raid_type": "raid_dp",
                    "parity_disk_count": 2,
                    "data_disk_count": 9,
                    "usable_size": 12309487,
                    "is_partition": true
                }
            ]
        }
    }
}

```

```

    },
    {
      "name": "node_1_SSD_1/plex1/rg0",
      "raid_type": "raid_dp",
      "parity_disk_count": 2,
      "data_disk_count": 9,
      "usable_size": 12309487,
      "is_partition": false
    }
  ]
},
"hybrid_cache": {
  "enabled": false
},
"mirror": {
  "enabled": true
}
},
]
"recommendation_spares" [
  {
    "node": {
      "uuid": "795bf7c2-fa4b-11e8-ba65-005056bbe5c1",
      "name": "node-2"
    }
  },
  "disk_class": "solid-state",
  "size": "2856845312",
  "is_partition": true,
  "disk_type": "ssd",
  "syncmirror_pool": "pool10",
  "usable": 1
}
],
"num_records": 2,
"warnings": [
  {
    "name": "node_1_SSD_1",
    "warning": {
      "code": 19726347,
      "message": "Unable to use all attached capacity on node \"node_1\".
3 local/remote pool disks not usable for mirroring.",
      "arguments": [
        "node_1",
        "3"
      ]
    }
  }
]

```



```

    ]
  },
  "action": {
    "code": 19726348,
    "message": "Contact technical support."
  }
},
{
  "name": "node_2_SSD_1",
  "warning": {
    "code": 19726347,
    "message": "Unable to use all attached capacity on node \"node_2\".
3 local/remote pool disks not usable for mirroring.",
    "arguments": [
      "node_2",
      "3"
    ]
  },
  "action": {
    "code": 19726348,
    "message": "Contact technical support."
  }
}
]
}

```

Retrieving the usable spare information for the cluster

The following example shows the response from retrieving usable spare information according to ONTAP best practices.

```

# The API:
/api/storage/aggregates?show_spares=true

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates?show_spares=true" -H
"accept: application/json"

# The response:
{
  "records": [],
  "num_records": 0,
  "spares": [
    {
      "node": {
        "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",

```

```

    "name": "node-2"
  },
  "disk_class": "solid_state",
  "disk_type": "ssd",
  "size": 3720609792,
  "checksum_style": "block",
  "syncmirror_pool": "pool0",
  "is_partitioned": true,
  "usable": 12,
  "layout_requirements": [
    {
      "raid_type": "raid_dp",
      "default": true,
      "aggregate_min_disks": 3,
      "raid_group": {
        "min": 3,
        "max": 28,
        "default": 24
      }
    },
    {
      "raid_type": "raid4",
      "default": false,
      "aggregate_min_disks": 2,
      "raid_group": {
        "min": 2,
        "max": 14,
        "default": 8
      }
    },
    {
      "raid_type": "raid_tec",
      "default": false,
      "aggregate_min_disks": 7,
      "raid_group": {
        "min": 4,
        "max": 29,
        "default": 25
      }
    }
  ]
},
{
  "node": {
    "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
    "name": "node-2"
  }
}

```

```

},
"disk_class": "solid_state",
"disk_type": "ssd_nvme",
"size": 8583380992,
"checksum_style": "block",
"syncmirror_pool": "pool0",
"is_partitioned": false,
"usable": 5,
"layout_requirements": [
  {
    "raid_type": "raid_dp",
    "default": true,
    "aggregate_min_disks": 3,
    "raid_group": {
      "min": 3,
      "max": 28,
      "default": 24
    }
  },
  {
    "raid_type": "raid4",
    "default": false,
    "aggregate_min_disks": 2,
    "raid_group": {
      "min": 2,
      "max": 14,
      "default": 8
    }
  },
  {
    "raid_type": "raid_tec",
    "default": false,
    "aggregate_min_disks": 7,
    "raid_group": {
      "min": 4,
      "max": 29,
      "default": 25
    }
  }
]
}
]
}

```

Retrieving the SSD spare count information for the cluster

The following example shows the response from retrieving SSD spare count information.

```
# The API:
/api/storage/aggregates?show_spares=true&flash_pool_eligible=true

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/aggregates?show_spares=true&flash_pool_eligible=true" -H
"accept: application/json"

# The response:
{
  "records": [],
  "num_records": 0,
  "spares": [
    {
      "node": {
        "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
        "name": "node-2"
      },
      "disk_class": "solid_state",
      "disk_type": "ssd",
      "size": 3720609792,
      "checksum_style": "block",
      "syncmirror_pool": "pool0",
      "is_partitioned": true,
      "usable": 12,
      "layout_requirements": [
        {
          "raid_type": "raid_dp",
          "default": true,
          "aggregate_min_disks": 3,
          "raid_group": {
            "min": 3,
            "max": 28,
            "default": 24
          }
        },
        {
          "raid_type": "raid4",
          "default": false,
          "aggregate_min_disks": 2,
          "raid_group": {
            "min": 2,
```

```

        "max": 14,
        "default": 8
    }
},
{
    "raid_type": "raid_tec",
    "default": false,
    "aggregate_min_disks": 7,
    "raid_group": {
        "min": 4,
        "max": 29,
        "default": 25
    }
}
]
},
{
    "node": {
        "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
        "name": "node-2"
    },
    "disk_class": "solid_state",
    "disk_type": "ssd_nvme",
    "size": 8583380992,
    "checksum_style": "block",
    "syncmirror_pool": "pool0",
    "is_partitioned": false,
    "usable": 5,
    "layout_requirements": [
        {
            "raid_type": "raid_dp",
            "default": true,
            "aggregate_min_disks": 3,
            "raid_group": {
                "min": 3,
                "max": 28,
                "default": 24
            }
        },
        {
            "raid_type": "raid4",
            "default": false,
            "aggregate_min_disks": 2,
            "raid_group": {
                "min": 2,
                "max": 14,

```

```

        "default": 8
    }
},
{
    "raid_type": "raid_tec",
    "default": false,
    "aggregate_min_disks": 7,
    "raid_group": {
        "min": 4,
        "max": 29,
        "default": 25
    }
}
]
}
]
}

```

Retrieving the total spare count information for the cluster

The following example shows the response from retrieving total spare count information, under advanced privilege.

```

# The API:
/api/storage/aggregates?show_spares=true&fields=**

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/aggregates?show_spares=true&fields=**" -H "accept:
application/json"

# The response:
{
  "records": [],
  "num_records": 0,
  "spares": [
    {
      "node": {
        "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
        "name": "node-2"
      },
      "disk_class": "solid_state",
      "disk_type": "ssd",
      "size": 3720609792,
      "checksum_style": "block",
      "syncmirror_pool": "pool0",

```

```
"is_partitioned": true,
"usable": 12,
"total": 14,
"layout_requirements": [
  {
    "raid_type": "raid_dp",
    "default": true,
    "aggregate_min_disks": 3,
    "raid_group": {
      "min": 3,
      "max": 28,
      "default": 24
    }
  },
  {
    "raid_type": "raid4",
    "default": false,
    "aggregate_min_disks": 2,
    "raid_group": {
      "min": 2,
      "max": 14,
      "default": 8
    }
  },
  {
    "raid_type": "raid_tec",
    "default": false,
    "aggregate_min_disks": 7,
    "raid_group": {
      "min": 4,
      "max": 29,
      "default": 25
    }
  }
],
{
  "node": {
    "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
    "name": "node-2"
  },
  "disk_class": "solid_state",
  "disk_type": "ssd_nvme",
  "size": 8583380992,
  "checksum_style": "block",
  "syncmirror_pool": "pool0",
```

```

"is_partitioned": false,
"usable": 5,
"total": 6,
"layout_requirements": [
  {
    "raid_type": "raid_dp",
    "default": true,
    "aggregate_min_disks": 3,
    "raid_group": {
      "min": 3,
      "max": 28,
      "default": 24
    }
  },
  {
    "raid_type": "raid4",
    "default": false,
    "aggregate_min_disks": 2,
    "raid_group": {
      "min": 2,
      "max": 14,
      "default": 8
    }
  },
  {
    "raid_type": "raid_tec",
    "default": false,
    "aggregate_min_disks": 7,
    "raid_group": {
      "min": 4,
      "max": 29,
      "default": 25
    }
  }
]
}
]
}

```

Simulating the creation of an aggregate using defined parameters

The following example shows the response in the simulation of a manual aggregate creation:

```

# The API:
/api/storage/aggregates

```



```

# The call
curl -X POST "https://<mgmt-ip>/api/storage/aggregates" -H "accept:
application/json" -d "{\"name\": \"node_2_SSD_1\", \"node\":
{\"name\": \"node-2\"}, \"block_storage\": {\"primary\": {\"disk_count\":
12\" }}, \"simulate\": \"true\"}"

# The response:
{
"records": [
  {
    "uuid": "795bf7c2-fa4b-11e8-ba65-005056bbe5c1",
    "name": "node_2_SSD_1",
    "node": {
      "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
      "name": "node-2",
    },
    "space": {
      "block_storage": {
        "size": 1116180480
      }
    },
    "block_storage": {
      "primary": {
        "disk_count": 12,
        "disk_class": "solid_state",
        "raid_type": "raid_dp",
        "disk_type": "ssd",
        "raid_size": 24,
        "simulated_raid_groups": [
          {
            "name": "node_2_SSD_1/plex0/rg0",
            "raid_type": "raid_dp",
            "parity_disk_count": 2,
            "data_disk_count": 10,
            "usable_size": 558090240,
            "is_partition": true
          },
        ]
      },
    },
    "hybrid_cache": {
      "enabled": false
    },
    "mirror": {
      "enabled": false
    }
  }
]
}

```

```

    }
  ],
  "num_records": 1,
  "warnings": [
    {
      "name": "node-2"
      "warning": {
        "code": 19726347,
        "message": "Number of unassigned disks attached to node \"node-2\":
6.",
        "arguments": [
          "node-2",
          "6"
        ]
      }
    }
  ]
}

```

Retrieve a collection of aggregates for an entire cluster

GET /storage/aggregates

Introduced In: 9.6

Retrieves the collection of aggregates for the entire cluster.

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.

- `metric.*`
- `space.block_storage.inactive_user_data`
- `space.block_storage.inactive_user_data_percent`
- `space.footprint`
- `is_spare_low`
- `statistics.*`

Related ONTAP commands

- `storage aggregate show`

Parameters

Name	Type	In	Required	Description
recommend	boolean	query	False	If set to 'true', it queries the system for the recommended optimal layout for creating new aggregates. The default setting is 'false'.
show_spares	boolean	query	False	If set to 'true', the spares object is returned instead of records to show the spare groups in the cluster. The default setting is 'false'.
flash_pool_eligible	boolean	query	False	If set to 'true' along with show_spares, the spares object is restricted to return spare groups that are compatible with whole disk based flash pool creation. The default setting is 'false'. <ul style="list-style-type: none"> Introduced in: 9.12
statistics.latency_read	integer	query	False	Filter by statistics.latency_read <ul style="list-style-type: none"> Introduced in: 9.7
statistics.latency_read.other	integer	query	False	Filter by statistics.latency_read.other <ul style="list-style-type: none"> Introduced in: 9.7

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

Name	Type	In	Required	Description
statistics.status	string	query	False	Filter by statistics.status • Introduced in: 9.7
statistics.throughput_raw.read	integer	query	False	Filter by statistics.throughput_raw.read • Introduced in: 9.7
statistics.throughput_raw.other	integer	query	False	Filter by statistics.throughput_raw.other • Introduced in: 9.7
statistics.throughput_raw.write	integer	query	False	Filter by statistics.throughput_raw.write • Introduced in: 9.7
statistics.throughput_raw.total	integer	query	False	Filter by statistics.throughput_raw.total • Introduced in: 9.7
data_encryption.drive_protection_enabled	boolean	query	False	Filter by data_encryption.drive_protection_enabled
data_encryption.software_encryption_enabled	boolean	query	False	Filter by data_encryption.software_encryption_enabled
sidl_enabled	boolean	query	False	Filter by sidel_enabled • Introduced in: 9.11

Name	Type	In	Required	Description
inactive_data_reporting.enabled	boolean	query	False	Filter by inactive_data_reporting.enabled • Introduced in: 9.8
inactive_data_reporting.start_time	string	query	False	Filter by inactive_data_reporting.start_time • Introduced in: 9.8
block_storage.storage_type	string	query	False	Filter by block_storage.storage_type • Introduced in: 9.11
block_storage.plexes.name	string	query	False	Filter by block_storage.plexes.name
block_storage.hybrid_cache.raid_size	integer	query	False	Filter by block_storage.hybrid_cache.raid_size • Introduced in: 9.12
block_storage.hybrid_cache.raid_type	string	query	False	Filter by block_storage.hybrid_cache.raid_type
block_storage.hybrid_cache.simulated_raid_groups.added_parity_disk_count	integer	query	False	Filter by block_storage.hybrid_cache.simulated_raid_groups.added_parity_disk_count • Introduced in: 9.12

Name	Type	In	Required	Description
block_storage.hybrid_cache.simulatedraid_groups.name	string	query	False	Filter by block_storage.hybrid_cache.simulatedraid_groups.name • Introduced in: 9.12
block_storage.hybrid_cache.simulatedraid_groups.existing_data_disk_count	integer	query	False	Filter by block_storage.hybrid_cache.simulatedraid_groups.existing_data_disk_count • Introduced in: 9.12
block_storage.hybrid_cache.simulatedraid_groups.usable_size	integer	query	False	Filter by block_storage.hybrid_cache.simulatedraid_groups.usable_size • Introduced in: 9.12
block_storage.hybrid_cache.simulatedraid_groups.existing_parity_disk_count	integer	query	False	Filter by block_storage.hybrid_cache.simulatedraid_groups.existing_parity_disk_count • Introduced in: 9.12
block_storage.hybrid_cache.simulatedraid_groups.added_data_disk_count	integer	query	False	Filter by block_storage.hybrid_cache.simulatedraid_groups.added_data_disk_count • Introduced in: 9.12

Name	Type	In	Required	Description
block_storage.hybrid_cache.simulatedraid_groups.is_partition	boolean	query	False	Filter by block_storage.hybrid_cache.simulatedraid_groups.is_partition • Introduced in: 9.12
block_storage.hybrid_cache.enabled	boolean	query	False	Filter by block_storage.hybrid_cache.enabled
block_storage.hybrid_cache.size	integer	query	False	Filter by block_storage.hybrid_cache.size
block_storage.hybrid_cache.storage_pools.storage_pool.uuid	string	query	False	Filter by block_storage.hybrid_cache.storage_pools.storage_pool.uuid • Introduced in: 9.11
block_storage.hybrid_cache.storage_pools.storage_pool.name	string	query	False	Filter by block_storage.hybrid_cache.storage_pools.storage_pool.name • Introduced in: 9.11
block_storage.hybrid_cache.storage_pools.allocation_units_count	integer	query	False	Filter by block_storage.hybrid_cache.storage_pools.allocation_units_count • Introduced in: 9.11
block_storage.hybrid_cache.used	integer	query	False	Filter by block_storage.hybrid_cache.used

Name	Type	In	Required	Description
block_storage.hybrid_cache.disk_count	integer	query	False	Filter by block_storage.hybrid_cache.disk_count
block_storage.hybrid_cache.disk_type	string	query	False	Filter by block_storage.hybrid_cache.disk_type • Introduced in: 9.12
block_storage.mirror.state	string	query	False	Filter by block_storage.mirror.state
block_storage.mirror.enabled	boolean	query	False	Filter by block_storage.mirror.enabled
block_storage.primary.disk_type	string	query	False	Filter by block_storage.primary.disk_type • Introduced in: 9.7
block_storage.primary.disk_class	string	query	False	Filter by block_storage.primary.disk_class
block_storage.primary.disk_count	integer	query	False	Filter by block_storage.primary.disk_count
block_storage.primary.simulated RAID groups.added_data_disk_count	integer	query	False	Filter by block_storage.primary.simulated RAID groups.added_data_disk_count • Introduced in: 9.11

Name	Type	In	Required	Description
block_storage.primary.simulated_raid_groups.is_partition	boolean	query	False	Filter by block_storage.primary.simulated_raid_groups.is_partition • Introduced in: 9.10
block_storage.primary.simulated_raid_groups.existing_parity_disk_count	integer	query	False	Filter by block_storage.primary.simulated_raid_groups.existing_parity_disk_count • Introduced in: 9.11
block_storage.primary.simulated_raid_groups.existing_data_disk_count	integer	query	False	Filter by block_storage.primary.simulated_raid_groups.existing_data_disk_count • Introduced in: 9.11
block_storage.primary.simulated_raid_groups.data_disk_count	integer	query	False	Filter by block_storage.primary.simulated_raid_groups.data_disk_count • Introduced in: 9.10
block_storage.primary.simulated_raid_groups.added_parity_disk_count	integer	query	False	Filter by block_storage.primary.simulated_raid_groups.added_parity_disk_count • Introduced in: 9.11

Name	Type	In	Required	Description
block_storage.primary.simulated_raid_groups.parity_disk_count	integer	query	False	Filter by block_storage.primary.simulated_raid_groups.parity_disk_count • Introduced in: 9.10
block_storage.primary.simulated_raid_groups.usable_size	integer	query	False	Filter by block_storage.primary.simulated_raid_groups.usable_size • Introduced in: 9.10
block_storage.primary.simulated_raid_groups.name	string	query	False	Filter by block_storage.primary.simulated_raid_groups.name • Introduced in: 9.10
block_storage.primary.simulated_raid_groups.raid_type	string	query	False	Filter by block_storage.primary.simulated_raid_groups.raid_type • Introduced in: 9.10
block_storage.primary.raid_type	string	query	False	Filter by block_storage.primary.raid_type
block_storage.primary.raid_size	integer	query	False	Filter by block_storage.primary.raid_size
block_storage.primary.checksum_style	string	query	False	Filter by block_storage.primary.checksum_style

Name	Type	In	Required	Description
block_storage.uses_partitions	boolean	query	False	Filter by block_storage.uses_partitions • Introduced in: 9.11
volume-count	integer	query	False	Filter by volume-count • Introduced in: 9.11
_tags	string	query	False	Filter by _tags • Introduced in: 9.13
metric.timestamp	string	query	False	Filter by metric.timestamp • Introduced in: 9.7
metric.status	string	query	False	Filter by metric.status • Introduced in: 9.7
metric.throughput.read	integer	query	False	Filter by metric.throughput.read • Introduced in: 9.7
metric.throughput.other	integer	query	False	Filter by metric.throughput.other • Introduced in: 9.7

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

Name	Type	In	Required	Description
metric.iops.other	integer	query	False	Filter by metric.iops.other • Introduced in: 9.7
metric.iops.write	integer	query	False	Filter by metric.iops.write • Introduced in: 9.7
metric.iops.total	integer	query	False	Filter by metric.iops.total • Introduced in: 9.7
home_node.uuid	string	query	False	Filter by home_node.uuid
home_node.name	string	query	False	Filter by home_node.name
dr_home_node.name	string	query	False	Filter by dr_home_node.name
dr_home_node.uuid	string	query	False	Filter by dr_home_node.uuid
recommendation_spaces.is_partition	boolean	query	False	Filter by recommendation_spaces.is_partition • Introduced in: 9.10
recommendation_spaces.disk_class	string	query	False	Filter by recommendation_spaces.disk_class • Introduced in: 9.10

Name	Type	In	Required	Description
recommendation_spares.disk_type	string	query	False	Filter by recommendation_spares.disk_type • Introduced in: 9.10
recommendation_spares.checksum_style	string	query	False	Filter by recommendation_spares.checksum_style • Introduced in: 9.10
recommendation_spares.syncmirror_pool	string	query	False	Filter by recommendation_spares.syncmirror_pool • Introduced in: 9.10
recommendation_spares.size	integer	query	False	Filter by recommendation_spares.size • Introduced in: 9.10
recommendation_spares.usable	integer	query	False	Filter by recommendation_spares.usable • Introduced in: 9.10
recommendation_spares.node.uuid	string	query	False	Filter by recommendation_spares.node.uuid • Introduced in: 9.10

Name	Type	In	Required	Description
recommendation_spare.node.name	string	query	False	Filter by recommendation_spare.node.name • Introduced in: 9.10
recommendation_spare.layout_requirements.raid_type	string	query	False	Filter by recommendation_spare.layout_requirements.raid_type • Introduced in: 9.10
recommendation_spare.layout_requirements.default	boolean	query	False	Filter by recommendation_spare.layout_requirements.default • Introduced in: 9.10
recommendation_spare.layout_requirements.raid_group.max	integer	query	False	Filter by recommendation_spare.layout_requirements.raid_group.max • Introduced in: 9.10
recommendation_spare.layout_requirements.raid_group.default	integer	query	False	Filter by recommendation_spare.layout_requirements.raid_group.default • Introduced in: 9.10
recommendation_spare.layout_requirements.raid_group.min	integer	query	False	Filter by recommendation_spare.layout_requirements.raid_group.min • Introduced in: 9.10

Name	Type	In	Required	Description
recommendation_spaces.layout_requirements.aggregate_min_disks	integer	query	False	Filter by recommendation_spaces.layout_requirements.aggregate_min_disks • Introduced in: 9.10
recommendation_spaces.total	integer	query	False	Filter by recommendation_spaces.total • Introduced in: 9.11
space.cloud_storage.used	integer	query	False	Filter by space.cloud_storage.used
space.footprint	integer	query	False	Filter by space.footprint
space.encyency.logical_used	integer	query	False	Filter by space.encyency.logical_used
space.encyency.cross_volume_inline_dedupe	boolean	query	False	Filter by space.encyency.cross_volume_inline_dedupe • Introduced in: 9.12
space.encyency.cross_volume_background_dedupe	boolean	query	False	Filter by space.encyency.cross_volume_background_dedupe • Introduced in: 9.12
space.encyency.ratio	number	query	False	Filter by space.encyency.ratio

Name	Type	In	Required	Description
space.encyency.savings	integer	query	False	Filter by space.encyency.savings
space.encyency.cross_volume_dedupe_savings	boolean	query	False	Filter by space.encyency.cross_volume_dedupe_savings • Introduced in: 9.12
space.encyency.auto_adaptive_compression_savings	boolean	query	False	Filter by space.encyency.auto_adaptive_compression_savings • Introduced in: 9.12
space.encyency_without_snapshots_flex_clones.logical_used	integer	query	False	Filter by space.encyency_without_snapshots_flex_clones.logical_used • Introduced in: 9.9
space.encyency_without_snapshots_flex_clones.savings	integer	query	False	Filter by space.encyency_without_snapshots_flex_clones.savings • Introduced in: 9.9
space.encyency_without_snapshots_flex_clones.ratio	number	query	False	Filter by space.encyency_without_snapshots_flex_clones.ratio • Introduced in: 9.9
space.block_storage.available	integer	query	False	Filter by space.block_storage.available

Name	Type	In	Required	Description
space.block_storage.size	integer	query	False	Filter by space.block_storage.size
space.block_storage.full_threshold_percent	integer	query	False	Filter by space.block_storage.full_threshold_percent
space.block_storage.volume_deduplication_space_saved	integer	query	False	Filter by space.block_storage.volume_deduplication_space_saved • Introduced in: 9.10
space.block_storage.volume_deduplication_space_saved_percent	integer	query	False	Filter by space.block_storage.volume_deduplication_space_saved_percent • Introduced in: 9.10
space.block_storage.data_compaction_space_saved	integer	query	False	Filter by space.block_storage.data_compaction_space_saved • Introduced in: 9.10
space.block_storage.data_compaction_space_saved_percent	integer	query	False	Filter by space.block_storage.data_compaction_space_saved_percent • Introduced in: 9.10

Name	Type	In	Required	Description
space.block_storage.data_compacted_count	integer	query	False	Filter by space.block_storage.data_compacted_count • Introduced in: 9.10
space.block_storage.used_including_snapshot_reserve_percent	integer	query	False	Filter by space.block_storage.used_including_snapshot_reserve_percent • Introduced in: 9.10
space.block_storage.used_including_snapshot_reserve	integer	query	False	Filter by space.block_storage.used_including_snapshot_reserve • 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: 9.10
space.block_storage.inactive_user_data	integer	query	False	Filter by space.block_storage.inactive_user_data
space.block_storage.volume_footprints_percent	integer	query	False	Filter by space.block_storage.volume_footprints_percent • Introduced in: 9.10
space.block_storage.used	integer	query	False	Filter by space.block_storage.used

Name	Type	In	Required	Description
space.block_storage.physical_used_percent	integer	query	False	Filter by space.block_storage.physical_used_percent • Introduced in: 9.10
space.block_storage.aggregate_metadata	integer	query	False	Filter by space.block_storage.aggregate_metadata • Introduced in: 9.10
space.block_storage.aggregate_metadata_percent	integer	query	False	Filter by space.block_storage.aggregate_metadata_percent • Introduced in: 9.10
space.block_storage.volume_deduplication_shared_count	integer	query	False	Filter by space.block_storage.volume_deduplication_shared_count • Introduced in: 9.10
space.block_storage.used_percent	integer	query	False	Filter by space.block_storage.used_percent • Introduced in: 9.13
space.block_storage.physical_used	integer	query	False	Filter by space.block_storage.physical_used • Introduced in: 9.9

Name	Type	In	Required	Description
space.snapshot.reset_percent	integer	query	False	Filter by space.snapshot.reset_percent • Introduced in: 9.10
space.snapshot.used_percent	integer	query	False	Filter by space.snapshot.used_percent • Introduced in: 9.10
space.snapshot.total	integer	query	False	Filter by space.snapshot.total • Introduced in: 9.10
space.snapshot.available	integer	query	False	Filter by space.snapshot.available • Introduced in: 9.10
space.snapshot.used	integer	query	False	Filter by space.snapshot.used • Introduced in: 9.10
space.ency_wit hout_snapshots.logical_used	integer	query	False	Filter by space.ency_wit hout_snapshots.logical_used
space.ency_wit hout_snapshots.savings	integer	query	False	Filter by space.ency_wit hout_snapshots.savings
space.ency_wit hout_snapshots.ratio	number	query	False	Filter by space.ency_wit hout_snapshots.ratio

Name	Type	In	Required	Description
node.uuid	string	query	False	Filter by node.uuid
node.name	string	query	False	Filter by node.name
snaplock_type	string	query	False	Filter by snaplock_type
name	string	query	False	Filter by name
create_time	string	query	False	Filter by create_time
state	string	query	False	Filter by state
uuid	string	query	False	Filter by uuid
inode_attributes.max_files_used	integer	query	False	Filter by inode_attributes.max_files_used • Introduced in: 9.11
inode_attributes.files_total	integer	query	False	Filter by inode_attributes.files_total • Introduced in: 9.11
inode_attributes.used_percent	integer	query	False	Filter by inode_attributes.used_percent • Introduced in: 9.11 • Max value: 100 • Min value: 0
inode_attributes.version	integer	query	False	Filter by inode_attributes.version • Introduced in: 9.11

Name	Type	In	Required	Description
inode_attributes.file_public_capacity	integer	query	False	Filter by inode_attributes.file_public_capacity • Introduced in: 9.11
inode_attributes.max_files_possible	integer	query	False	Filter by inode_attributes.max_files_possible • Introduced in: 9.11
inode_attributes.max_files_available	integer	query	False	Filter by inode_attributes.max_files_available • Introduced in: 9.11
inode_attributes.file_private_capacity	integer	query	False	Filter by inode_attributes.file_private_capacity • Introduced in: 9.11
inode_attributes.files_private_used	integer	query	False	Filter by inode_attributes.files_private_used • Introduced in: 9.11
inode_attributes.files_used	integer	query	False	Filter by inode_attributes.files_used • Introduced in: 9.11
is_spare_low	boolean	query	False	Filter by is_spare_low • Introduced in: 9.11

Name	Type	In	Required	Description
snapshot.max_files_available	integer	query	False	Filter by snapshot.max_files_available • Introduced in: 9.10
snapshot.files_total	integer	query	False	Filter by snapshot.files_total • Introduced in: 9.10
snapshot.max_files_used	integer	query	False	Filter by snapshot.max_files_used • Introduced in: 9.10
snapshot.files_used	integer	query	False	Filter by snapshot.files_used • 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.
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

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

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
error	error	
num_records	integer	Number of records
records	array[aggregate]	
spares	array[aggregate_spare]	
warnings	array[aggregate_warning]	List of warnings and remediation advice for the aggregate recommendation.

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "_tags": [
      "team:csi",
      "environment:test"
    ],
    "block_storage": {
      "hybrid_cache": {
        "disk_count": 6,
        "disk_type": "fc",
        "raid_size": 24,
        "raid_type": "raid_dp",
        "simulated_raid_groups": {
        },
        "size": 1612709888,
        "storage_pools": {
          "storage_pool": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            }
          }
        }
      }
    }
  }
}
```

```

    },
    "name": "storage_pool_1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"used": 26501122
},
"mirror": {
  "enabled": "",
  "state": "unmirrored"
},
"plexes": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "plex0"
},
"primary": {
  "checksum_style": "block",
  "disk_class": "performance",
  "disk_count": 8,
  "disk_type": "fc",
  "raid_size": 16,
  "raid_type": "raid_dp",
  "simulated_raid_groups": {
    "raid_type": "raid_dp"
  }
},
"storage_type": "hdd"
},
"cloud_storage": {
  "stores": {
    "cloud_store": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "store1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "used": 0
  }
},
},

```

```
"create_time": "2018-01-01 16:00:00 +0000",
"dr_home_node": {
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"home_node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"inactive_data_reporting": {
  "start_time": "2019-12-12 16:00:00 +0000"
},
"inode_attributes": {
  "file_private_capacity": 31136,
  "file_public_capacity": 31136,
  "files_private_used": 502,
  "files_total": 31136,
  "files_used": 97,
  "max_files_available": 31136,
  "max_files_possible": 2844525,
  "max_files_used": 97,
  "used_percent": 5,
  "version": 4
},
"is_spare_low": "",
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"duration": "PT15S",
"iops": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"latency": {
  "read": 200,
  "total": 1000,
  "write": 100
}
```

```

    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "name": "node1_aggr_1",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "recommendation_spares": {
    "checksum_style": "block",
    "disk_class": "solid_state",
    "disk_type": "fc",
    "is_partition": 1,
    "layout_requirements": {
      "aggregate_min_disks": 6,
      "raid_group": {
        "default": 16,
        "max": 28,
        "min": 5
      },
      "raid_type": "raid_dp"
    },
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "size": 10156769280,
    "syncmirror_pool": "pool0",
    "total": 10,
    "usable": 9
  }
}

```

```

},
"snaplock_type": "non_snaplock",
"snapshot": {
  "files_total": 10,
  "files_used": 3,
  "max_files_available": 5,
  "max_files_used": 50
},
"space": {
  "block_storage": {
    "aggregate_metadata": 2655,
    "aggregate_metadata_percent": 8,
    "available": 10156560384,
    "data_compacted_count": 1990000,
    "data_compaction_space_saved": 1996000,
    "data_compaction_space_saved_percent": 27,
    "full_threshold_percent": 0,
    "inactive_user_data": 304448,
    "inactive_user_data_percent": 0,
    "physical_used": 2461696,
    "physical_used_percent": 50,
    "size": 10156769280,
    "used": 2088960,
    "used_including_snapshot_reserve": 674685,
    "used_including_snapshot_reserve_percent": 35,
    "used_percent": 50,
    "volume_deduplication_shared_count": 1990000,
    "volume_deduplication_space_saved": 1996000,
    "volume_deduplication_space_saved_percent": 27,
    "volume_footprints_percent": 14
  },
  "cloud_storage": {
    "used": 402743264
  },
  "efficiency": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots_flexclones": {
    "logical_used": 0,

```

```

    "ratio": 0,
    "savings": 0
  },
  "footprint": 608896,
  "snapshot": {
    "available": 2000,
    "reserve_percent": 20,
    "total": 5000,
    "used": 3000,
    "used_percent": 45
  }
},
"state": "online",
"statistics": {
  "iops_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"uuid": "string",
"volume-count": 0
},
"spares": {
  "checksum_style": "block",
  "disk_class": "solid_state",
  "disk_type": "fc",
  "is_partition": 1,
  "layout_requirements": {
    "aggregate_min_disks": 6,
    "raid_group": {
      "default": 16,
      "max": 28,
      "min": 5
    }
  }
}

```



```

    },
    "raid_type": "raid_dp"
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "size": 10156769280,
  "syncmirror_pool": "pool0",
  "total": 10,
  "usable": 9
},
"warnings": {
  "action": {
    "arguments": {
    }
  },
  "warning": {
    "arguments": {
    }
  }
}
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
787092	The target field cannot be specified for this operation.
918138	Internal error. Failed to get encryption operation status.
8586225	Encountered unexpected error in retrieving metrics and statistics for an aggregate.
19726341	Not enough eligible spare disks are available on the node.

Error Code	Description
19726344	No recommendation can be made for this cluster.
19726357	Aggregate recommendations are not supported on MetroCluster with Fibre Channel (FC).
19726358	Aggregate recommendations are not supported on ONTAP Cloud.
19726382	Another provisioning operation is in progress on this cluster. Wait a few minutes, and try the operation again.
19726386	Encountered an error when retrieving licensing information on this cluster.
19726387	No recommendation can be provided for this cluster within the license capacity.
19726401	Aggregate recommendations are not supported when the DR group is not in the "normal" state.
19726402	Internal error. Unable to determine the MetroCluster configuration state.
19726403	Aggregate recommendation is not supported when there are no healthy target connections to remote storage.
19726404	The recommended mirrored aggregate couldn't use all the attached capacity in one of the SyncMirror pools. Make sure that the remote and local storage is symmetrically wired.
19726405	Not all local and remote disks attached to the node have been auto-partitioned.
19726406	Aggregate recommendations are not supported on this node because remote and local storage is not symmetrically wired.
19726540	The next tag is not supported for recommended aggregates. Retry the operation with a higher "return_timeout" value.
196608055	Aggregate recommendation is not supported on this node because it does not support NetApp Aggregate Encryption (NAE).
196608206	Internal error. Failed to get encryption operation status.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

_links

Name	Type	Description
self	href	

simulated_raid_groups

Name	Type	Description
added_data_disk_count	integer	Number of added data disks in RAID group.

Name	Type	Description
added_parity_disk_count	integer	Number of added parity disks in RAID group.
existing_data_disk_count	integer	Number of existing data disks in the RAID group.
existing_parity_disk_count	integer	Number of existing parity disks in the RAID group.
is_partition	boolean	Indicates whether the disk is partitioned (true) or whole (false).
name	string	Name of the raid group.
usable_size	integer	Usable size of each disk, in bytes.

storage_pool_reference

Shared Storage Pool

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pools

Name	Type	Description
allocation_units_count	integer	Allocation count of storage pool.
storage_pool	storage_pool_reference	Shared Storage Pool

hybrid_cache

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

Name	Type	Description
disk_count	integer	Number of disks used in the cache tier of the aggregate. Only provided when hybrid_cache.enabled is 'true'.

Name	Type	Description
disk_type	string	Type of disk being used by the aggregate's cache tier.
enabled	boolean	Specifies whether the aggregate uses HDDs with SSDs as a cache.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
simulated_raid_groups	array[simulated_raid_groups]	
size	integer	Total usable space in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.
storage_pools	array[storage_pools]	List of storage pool properties and allocation_units_count for aggregate.
used	integer	Space used in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex_reference

Plex

Name	Type	Description
_links	_links	
name	string	

simulated_raid_groups

Name	Type	Description
added_data_disk_count	integer	Number of added data disks in RAID group.
added_parity_disk_count	integer	Number of added parity disks in RAID group.
data_disk_count	integer	Number of data disks in RAID group.
existing_data_disk_count	integer	Number of existing data disks in the RAID group.
existing_parity_disk_count	integer	Number of existing parity disks in the RAID group.
is_partition	boolean	Indicates whether the disk is partitioned (true) or whole (false).
name	string	Name of the raid group.
parity_disk_count	integer	Number of parity disks in RAID group.
raid_type	string	RAID type of the aggregate.
usable_size	integer	Usable size of each disk, in bytes.

primary

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

Name	Type	Description
checksum_style	string	The checksum style used by the aggregate.
disk_class	string	The class of disks being used by the aggregate.
disk_count	integer	Number of disks used in the aggregate. This includes parity disks, but excludes disks in the hybrid cache.

Name	Type	Description
disk_type	string	The type of disk being used by the aggregate.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type of the aggregate.
simulated_raid_groups	array[simulated_raid_groups]	

block_storage

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

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

cloud_store

Cloud store

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cloud_storage_tier

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

cloud_storage

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

Name	Type	Description
attach_eligible	boolean	Specifies whether the aggregate is eligible for a cloud store to be attached.
migrate_threshold	integer	Specifies the minimum percentage of performance tier free space that must exist in order for migration of data from the capacity tier to performance tier to be allowed. Only valid for PATCH operations.
stores	array[cloud_storage_tier]	Configuration information for each cloud storage portion of the aggregate.
tiering_fullness_threshold	integer	The percentage of space in the performance tier that must be used before data is tiered out to the cloud store. Only valid for PATCH operations.

data_encryption

Name	Type	Description
drive_protection_enabled	boolean	Specifies whether the aggregate uses self-encrypting drives with data protection enabled.

Name	Type	Description
software_encryption_enabled	boolean	Specifies whether NetApp aggregate encryption is enabled. All data in the aggregate is encrypted.

dr_home_node

Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.

Name	Type	Description
name	string	
uuid	string	

home_node

Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

inactive_data_reporting

Name	Type	Description
enabled	boolean	Specifies whether or not inactive data reporting is enabled on the aggregate.
start_time	string	Timestamp at which inactive data reporting was enabled on the aggregate.

inode_attributes

Name	Type	Description
file_private_capacity	integer	Number of files that can currently be stored on disk for system metadata files. This number will dynamically increase as more system files are created. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .
file_public_capacity	integer	Number of files that can currently be stored on disk for user-visible files. This number will dynamically increase as more user-visible files are created. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .
files_private_used	integer	Number of system metadata files used. If the referenced file system is restricted or offline, a value of 0 is returned. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .
files_total	integer	Maximum number of user-visible files that this referenced file system can currently hold. If the referenced file system is restricted or offline, a value of 0 is returned.

Name	Type	Description
files_used	integer	Number of user-visible files used in the referenced file system. If the referenced file system is restricted or offline, a value of 0 is returned.
max_files_available	integer	The count of the maximum number of user-visible files currently allowable on the referenced file system.
max_files_possible	integer	The largest value to which the maxfiles-available parameter can be increased by reconfiguration, on the referenced file system.
max_files_used	integer	The number of user-visible files currently in use on the referenced file system.
used_percent	integer	The percentage of disk space currently in use based on user-visible file count on the referenced file system.
version	integer	The inofile-version of the aggregate. If the referenced file system is restricted or offline, a value of 0 is returned. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .

iops

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

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

latency

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

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

throughput

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

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

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

metric

The most recent sample of I/O metrics for the aggregate.

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

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.

node

Node where the aggregate currently resides.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

raid_group

Name	Type	Description
default	integer	Default number of disks in a RAID group.

Name	Type	Description
max	integer	Maximum number of disks allowed in a RAID group.
min	integer	Minimum number of disks allowed in a RAID group.

layout_requirement

Name	Type	Description
aggregate_min_disks	integer	Minimum number of disks to create an aggregate.
default	boolean	Indicates if this RAID type is the default.
raid_group	raid_group	
raid_type	string	RAID type.

node

Node where the spares are assigned.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

aggregate_spare

Name	Type	Description
checksum_style	string	The checksum type that has been assigned to the spares.
disk_class	string	Disk class of spares.
disk_type	string	Type of disk.
is_partition	boolean	Indicates whether a disk is partitioned (true) or whole (false)
layout_requirements	array[layout_requirement]	Available RAID protections and their restrictions.

Name	Type	Description
node	node	Node where the spares are assigned.
size	integer	Usable size of each spare, in bytes.
syncmirror_pool	string	SyncMirror spare pool.
total	integer	Total number of spares in the bucket. The total spare count for each class of spares also includes reserved spare capacity recommended by ONTAP best practices. <ul style="list-style-type: none"> • example: 10 • readOnly: 1 • Introduced in: 9.11 • x-nullable: true
usable	integer	Total number of usable spares in the bucket. The usable count for each class of spares does not include reserved spare capacity recommended by ONTAP best practices. <ul style="list-style-type: none"> • example: 9 • readOnly: 1 • Introduced in: 9.6 • x-nullable: true

snapshot

Name	Type	Description
files_total	integer	Total files allowed in Snapshot copies
files_used	integer	Total files created in Snapshot copies
max_files_available	integer	Maximum files available for Snapshot copies

Name	Type	Description
max_files_used	integer	Files in use by Snapshot copies

block_storage

Name	Type	Description
aggregate_metadata	integer	Space used by different metafiles and internal operations inside the aggregate, in bytes.
aggregate_metadata_percent	integer	Aggregate metadata as a percentage.
available	integer	Space available in bytes.
data_compacted_count	integer	Amount of compacted data in bytes.
data_compaction_space_saved	integer	Space saved in bytes by compacting the data.
data_compaction_space_saved_percent	integer	Percentage saved by compacting the data.
full_threshold_percent	integer	The aggregate used percentage at which 'monitor.volume.full' EMS is generated.
inactive_user_data	integer	The size that is physically used in the block storage and has a cold temperature, in bytes. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>block_storage.inactive_user_data</code> or <code>**</code> .

Name	Type	Description
inactive_user_data_percent	integer	The percentage of inactive user data in the block storage. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>block_storage.inactive_user_data_percent</code> or <code>**</code> .
physical_used	integer	Total physical used size of an aggregate in bytes.
physical_used_percent	integer	Physical used percentage.
size	integer	Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
used	integer	Space used or reserved in bytes. Includes volume guarantees and aggregate metadata.
used_including_snapshot_reserve	integer	Total used including the Snapshot copy reserve, in bytes.
used_including_snapshot_reserve_percent	integer	Total used including the Snapshot reserve as a percentage.
used_percent	integer	Aggregate used percentage.
volume_deduplication_shared_count	integer	Amount of shared bytes counted by storage efficiency.
volume_deduplication_space_saved	integer	Amount of space saved in bytes by storage efficiency.
volume_deduplication_space_saved_percent	integer	Percentage of space saved by storage efficiency.

Name	Type	Description
volume_footprints_percent	integer	A summation of volume footprints inside the aggregate, as a percentage. A volume's footprint is the amount of space being used for the volume in the aggregate.

cloud_storage

Name	Type	Description
used	integer	Used space in bytes in the cloud store. Only applicable for aggregates with a cloud store tier.

efficiency

Storage efficiency.

Name	Type	Description
auto_adaptive_compression_savings	boolean	Indicates whether or not aggregate has auto adaptive compression savings.
cross_volume_background_dedupe	boolean	Indicates whether or not cross volume background deduplication is enabled.
cross_volume_dedupe_savings	boolean	Indicates whether or not aggregate has cross volume deduplication savings.
cross_volume_inline_dedupe	boolean	Indicates whether or not cross volume inline deduplication is enabled.
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

efficiency_without_snapshots

Storage efficiency that does not include the savings provided by Snapshot copies.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

efficiency_without_snapshots_flexclones

Storage efficiency that does not include the savings provided by Snapshot copies and Flexclone volumes.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

snapshot

Name	Type	Description
available	integer	Available space for Snapshot copies in bytes
reserve_percent	integer	Percentage of space reserved for Snapshot copies
total	integer	Total space for Snapshot copies in bytes
used	integer	Space used by Snapshot copies in bytes
used_percent	integer	Percentage of disk space used by Snapshot copies

space

Name	Type	Description
block_storage	block_storage	

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

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.

Name	Type	Description
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.

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

The real time I/O statistics for the aggregate.

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.
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.

aggregate

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.
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation.
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	home_node	Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.
inactive_data_reporting	inactive_data_reporting	
inode_attributes	inode_attributes	
is_spare_low	boolean	Specifies whether the aggregate is in a spares low condition on any of the RAID groups. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .

Name	Type	Description
metric	metric	The most recent sample of I/O metrics for the aggregate.
name	string	Aggregate name.
node	node	Node where the aggregate currently resides.
recommendation_spares	array[aggregate_spare]	Information on the aggregate's remaining hot spare disks.
sidl_enabled	boolean	Specifies whether or not SIDL is enabled on the aggregate.
snaplock_type	string	SnapLock type.
snapshot	snapshot	
space	space	
state	string	Operational state of the aggregate.
statistics	statistics	The real time I/O statistics for the aggregate.
uuid	string	Aggregate UUID.
volume-count	integer	Number of volumes in the aggregate.

action

Name	Type	Description
arguments	array[string]	Arguments present in the specified action message.
code	integer	Corrective action code of the specified action.
message	string	Specifies the corrective action to be taken to resolve the issue.

warning

Name	Type	Description
arguments	array[string]	Arguments present in the warning message encountered.
code	integer	Warning code of the warning encountered.
message	string	Details of the warning encountered by the aggregate simulate query.

aggregate_warning

Name	Type	Description
action	action	
name	string	Name of the entity that returns the warning.
warning	warning	

Create a collection of aggregates for an entire cluster

POST /storage/aggregates

Introduced In: 9.6

Automatically creates aggregates based on an optimal layout recommended by the system. Alternatively, properties can be provided to create an aggregate according to the requested specification. This request starts a job and returns a link to that job. POST operations will be blocked while one or more nodes in the cluster are simulating or implementing automatic aggregate creation.

Required properties

Properties are not required for this API. The following properties are only required if you want to specify properties for aggregate creation:

- name - Name of the aggregate.
- node.name or node.uuid - Node on which the aggregate will be created.
- block_storage.primary.disk_count - Number of disks to be used to create the aggregate.

Default values

If not specified in POST, the following default values are assigned. The remaining unspecified properties will receive system dependent default values.

- block_storage.mirror.enabled - *false*

- snaplock_type - *non_snaplock*

Related ONTAP commands

- storage aggregate auto-provision
- storage aggregate create

Example:

```
POST /api/storage/aggregates {"node": {"name": "node1"}, "name": "test",  
"block_storage": {"primary": {"disk_count": "10"}}}
```

Parameters

Name	Type	In	Required	Description
disk_size	integer	query	False	If set, POST only selects disks of the specified size.

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

Request Body

Name	Type	Description
_links	_links	
_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
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation.
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	home_node	Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.
inactive_data_reporting	inactive_data_reporting	
inode_attributes	inode_attributes	
is_spare_low	boolean	Specifies whether the aggregate is in a spares low condition on any of the RAID groups. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or **.
metric	metric	The most recent sample of I/O metrics for the aggregate.
name	string	Aggregate name.
node	node	Node where the aggregate currently resides.

Name	Type	Description
recommendation_spares	array[aggregate_spare]	Information on the aggregate's remaining hot spare disks.
sidl_enabled	boolean	Specifies whether or not SIDL is enabled on the aggregate.
snaplock_type	string	SnapLock type.
snapshot	snapshot	
space	space	
state	string	Operational state of the aggregate.
statistics	statistics	The real time I/O statistics for the aggregate.
uuid	string	Aggregate UUID.
volume-count	integer	Number of volumes in the aggregate.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "disk_type": "fc",
      "raid_size": 24,
      "raid_type": "raid_dp",
      "simulated_raid_groups": {
      },
      "size": 1612709888,
      "storage_pools": {
        "storage_pool": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "storage_pool_1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "used": 26501122
    },
    "mirror": {
      "enabled": "",
      "state": "unmirrored"
    },
    "plexes": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "plex0"
    }
  },
}
```



```

"primary": {
  "checksum_style": "block",
  "disk_class": "performance",
  "disk_count": 8,
  "disk_type": "fc",
  "raid_size": 16,
  "raid_type": "raid_dp",
  "simulated_raid_groups": {
    "raid_type": "raid_dp"
  }
},
"storage_type": "hdd"
},
"cloud_storage": {
  "stores": {
    "cloud_store": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "store1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "used": 0
  }
},
"create_time": "2018-01-01 16:00:00 +0000",
"dr_home_node": {
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"home_node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"inactive_data_reporting": {
  "start_time": "2019-12-12 16:00:00 +0000"
},
"inode_attributes": {
  "file_private_capacity": 31136,

```

```

    "file_public_capacity": 31136,
    "files_private_used": 502,
    "files_total": 31136,
    "files_used": 97,
    "max_files_available": 31136,
    "max_files_possible": 2844525,
    "max_files_used": 97,
    "used_percent": 5,
    "version": 4
  },
  "is_spare_low": "",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"name": "node1",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},

```

```

"recommendation_spares": {
  "checksum_style": "block",
  "disk_class": "solid_state",
  "disk_type": "fc",
  "is_partition": 1,
  "layout_requirements": {
    "aggregate_min_disks": 6,
    "raid_group": {
      "default": 16,
      "max": 28,
      "min": 5
    },
    "raid_type": "raid_dp"
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "size": 10156769280,
  "syncmirror_pool": "pool0",
  "total": 10,
  "usable": 9
},
"snaplock_type": "non_snaplock",
"snapshot": {
  "files_total": 10,
  "files_used": 3,
  "max_files_available": 5,
  "max_files_used": 50
},
"space": {
  "block_storage": {
    "aggregate_metadata": 2655,
    "aggregate_metadata_percent": 8,
    "available": 10156560384,
    "data_compacted_count": 1990000,
    "data_compaction_space_saved": 1996000,
    "data_compaction_space_saved_percent": 27,
    "full_threshold_percent": 0,
    "inactive_user_data": 304448,
    "inactive_user_data_percent": 0,

```

```

    "physical_used": 2461696,
    "physical_used_percent": 50,
    "size": 10156769280,
    "used": 2088960,
    "used_including_snapshot_reserve": 674685,
    "used_including_snapshot_reserve_percent": 35,
    "used_percent": 50,
    "volume_deduplication_shared_count": 1990000,
    "volume_deduplication_space_saved": 1996000,
    "volume_deduplication_space_saved_percent": 27,
    "volume_footprints_percent": 14
  },
  "cloud_storage": {
    "used": 402743264
  },
  "efficiency": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots_flexclones": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "footprint": 608896,
  "snapshot": {
    "available": 2000,
    "reserve_percent": 20,
    "total": 5000,
    "used": 3000,
    "used_percent": 45
  }
},
"state": "online",
"statistics": {
  "iops_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  }
},

```

```

"latency_raw": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"status": "ok",
"throughput_raw": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 +0000"
},
"uuid": "string",
"volume-count": 0
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	
num_records	integer	Number of records
records	array[aggregate]	
warnings	array[aggregate_warning]	List of validation warnings and remediation advice for the aggregate simulate behavior.

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "disk_type": "fc",
      "raid_size": 24,
      "raid_type": "raid_dp",
      "simulated_raid_groups": {
      },
      "size": 1612709888,
      "storage_pools": {
        "storage_pool": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          }
        },
        "name": "storage_pool_1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "used": 26501122
  },
  "mirror": {
    "enabled": "",
  }
}
```

```

    "state": "unmirrored"
  },
  "plexes": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "plex0"
  },
  "primary": {
    "checksum_style": "block",
    "disk_class": "performance",
    "disk_count": 8,
    "disk_type": "fc",
    "raid_size": 16,
    "raid_type": "raid_dp",
    "simulated_raid_groups": {
      "raid_type": "raid_dp"
    }
  },
  "storage_type": "hdd"
},
"cloud_storage": {
  "stores": {
    "cloud_store": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "store1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "used": 0
  }
},
"create_time": "2018-01-01 16:00:00 +0000",
"dr_home_node": {
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"home_node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
}

```

```

    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"inactive_data_reporting": {
  "start_time": "2019-12-12 16:00:00 +0000"
},
"inode_attributes": {
  "file_private_capacity": 31136,
  "file_public_capacity": 31136,
  "files_private_used": 502,
  "files_total": 31136,
  "files_used": 97,
  "max_files_available": 31136,
  "max_files_possible": 2844525,
  "max_files_used": 97,
  "used_percent": 5,
  "version": 4
},
"is_spare_low": "",
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"duration": "PT15S",
"iops": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"latency": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"status": "ok",
"throughput": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 +0000"
},

```



```
"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"recommendation_spare": {
  "checksum_style": "block",
  "disk_class": "solid_state",
  "disk_type": "fc",
  "is_partition": 1,
  "layout_requirements": {
    "aggregate_min_disks": 6,
    "raid_group": {
      "default": 16,
      "max": 28,
      "min": 5
    }
  },
  "raid_type": "raid_dp"
},
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": 10156769280,
"syncmirror_pool": "pool0",
"total": 10,
"usable": 9
},
"snaplock_type": "non_snaplock",
"snapshot": {
  "files_total": 10,
  "files_used": 3,
  "max_files_available": 5,
  "max_files_used": 50
},
"space": {
```

```

"block_storage": {
  "aggregate_metadata": 2655,
  "aggregate_metadata_percent": 8,
  "available": 10156560384,
  "data_compacted_count": 1990000,
  "data_compaction_space_saved": 1996000,
  "data_compaction_space_saved_percent": 27,
  "full_threshold_percent": 0,
  "inactive_user_data": 304448,
  "inactive_user_data_percent": 0,
  "physical_used": 2461696,
  "physical_used_percent": 50,
  "size": 10156769280,
  "used": 2088960,
  "used_including_snapshot_reserve": 674685,
  "used_including_snapshot_reserve_percent": 35,
  "used_percent": 50,
  "volume_deduplication_shared_count": 1990000,
  "volume_deduplication_space_saved": 1996000,
  "volume_deduplication_space_saved_percent": 27,
  "volume_footprints_percent": 14
},
"cloud_storage": {
  "used": 402743264
},
"efficiency": {
  "logical_used": 0,
  "ratio": 0,
  "savings": 0
},
"efficiency_without_snapshots": {
  "logical_used": 0,
  "ratio": 0,
  "savings": 0
},
"efficiency_without_snapshots_flexclones": {
  "logical_used": 0,
  "ratio": 0,
  "savings": 0
},
"footprint": 608896,
"snapshot": {
  "available": 2000,
  "reserve_percent": 20,
  "total": 5000,
  "used": 3000,

```

```

    "used_percent": 45
  }
},
"state": "online",
"statistics": {
  "iops_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"uuid": "string",
"volume-count": 0
},
"warnings": {
  "action": {
    "arguments": {
    }
  },
  "warning": {
    "arguments": {
    }
  }
}
}
}

```

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	The aggregate create job failed to create the aggregate.
786438	Failed to create an aggregate on the node.
786439	An aggregate already uses the specified name.
786446	The node is not in cluster.
786468	VLDB is offline.
786819	The value is invalid for the specified option at the current privilege level.
786902	RAID-TEC aggregate is not fully supported.
786911	Not every node in the cluster has the Data ONTAP version required for the feature.
787069	Node is setup for MetroCluster over IP configuration; creating an unmirrored aggregate is not supported in this configuration.
787092	The target field cannot be specified for this operation.
918138	Internal error. Failed to get encryption operation status.
1114292	The required SnapLock license is not installed.
2425736	No matching node found for the target UUID.
19726341	Not enough eligible spare disks are available on the node.
19726344	No recommendation can be made for this cluster.
19726357	Automatic aggregate creation is not supported on MetroCluster with Fibre Channel (FC).
19726358	Automatic aggregate creation is not supported on ONTAP Cloud.
19726373	Recommendation specified for creating aggregates is not current.

Error Code	Description
19726378	Failed to create recommended aggregates on one or more nodes.
19726382	Another provisioning operation is in progress on this cluster. Wait a few minutes, and try the operation again.
19726386	Encountered an error when retrieving licensing information on this cluster.
19726387	No recommendation can be provided for this cluster within the license capacity.
19726401	Aggregate recommendations are not supported when the DR group is not in the "normal" state.
19726402	Internal error. Unable to determine the MetroCluster configuration state.
19726403	Aggregate recommendation is not supported when there are no healthy target connections to remote storage.
196608055	Aggregate recommendation is not supported on this node because it does not support NetApp Aggregate Encryption (NAE).
196608206	Internal error. Failed to get encryption operation status.

Name	Type	Description
error	error	

Example error

```

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

simulated_raid_groups

Name	Type	Description
added_data_disk_count	integer	Number of added data disks in RAID group.
added_parity_disk_count	integer	Number of added parity disks in RAID group.
existing_data_disk_count	integer	Number of existing data disks in the RAID group.
existing_parity_disk_count	integer	Number of existing parity disks in the RAID group.
is_partition	boolean	Indicates whether the disk is partitioned (true) or whole (false).
name	string	Name of the raid group.
usable_size	integer	Usable size of each disk, in bytes.

storage_pool_reference

Shared Storage Pool

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pools

Name	Type	Description
allocation_units_count	integer	Allocation count of storage pool.
storage_pool	storage_pool_reference	Shared Storage Pool

hybrid_cache

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

Name	Type	Description
disk_count	integer	Number of disks used in the cache tier of the aggregate. Only provided when <code>hybrid_cache.enabled</code> is 'true'.
disk_type	string	Type of disk being used by the aggregate's cache tier.
enabled	boolean	Specifies whether the aggregate uses HDDs with SSDs as a cache.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when <code>hybrid_cache.enabled</code> is 'true'.
simulated_raid_groups	array[simulated_raid_groups]	
size	integer	Total usable space in bytes of SSD cache. Only provided when <code>hybrid_cache.enabled</code> is 'true'.
storage_pools	array[storage_pools]	List of storage pool properties and <code>allocation_units_count</code> for aggregate.
used	integer	Space used in bytes of SSD cache. Only provided when <code>hybrid_cache.enabled</code> is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex_reference

Plex

Name	Type	Description
_links	_links	
name	string	

simulated_raid_groups

Name	Type	Description
added_data_disk_count	integer	Number of added data disks in RAID group.
added_parity_disk_count	integer	Number of added parity disks in RAID group.
data_disk_count	integer	Number of data disks in RAID group.
existing_data_disk_count	integer	Number of existing data disks in the RAID group.
existing_parity_disk_count	integer	Number of existing parity disks in the RAID group.
is_partition	boolean	Indicates whether the disk is partitioned (true) or whole (false).
name	string	Name of the raid group.
parity_disk_count	integer	Number of parity disks in RAID group.
raid_type	string	RAID type of the aggregate.
usable_size	integer	Usable size of each disk, in bytes.

primary

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

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

block_storage

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

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

Name	Type	Description
uses_partitions	boolean	If true, aggregate is using shared disks.

cloud_store

Cloud store

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cloud_storage_tier

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

cloud_storage

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

Name	Type	Description
attach_eligible	boolean	Specifies whether the aggregate is eligible for a cloud store to be attached.
migrate_threshold	integer	Specifies the minimum percentage of performance tier free space that must exist in order for migration of data from the capacity tier to performance tier to be allowed. Only valid for PATCH operations.
stores	array[cloud_storage_tier]	Configuration information for each cloud storage portion of the aggregate.

Name	Type	Description
tiering_fullness_threshold	integer	The percentage of space in the performance tier that must be used before data is tiered out to the cloud store. Only valid for PATCH operations.

data_encryption

Name	Type	Description
drive_protection_enabled	boolean	Specifies whether the aggregate uses self-encrypting drives with data protection enabled.
software_encryption_enabled	boolean	Specifies whether NetApp aggregate encryption is enabled. All data in the aggregate is encrypted.

dr_home_node

Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.

Name	Type	Description
name	string	
uuid	string	

home_node

Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

inactive_data_reporting

Name	Type	Description
enabled	boolean	Specifies whether or not inactive data reporting is enabled on the aggregate.

Name	Type	Description
start_time	string	Timestamp at which inactive data reporting was enabled on the aggregate.

inode_attributes

Name	Type	Description
file_private_capacity	integer	Number of files that can currently be stored on disk for system metadata files. This number will dynamically increase as more system files are created. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .
file_public_capacity	integer	Number of files that can currently be stored on disk for user-visible files. This number will dynamically increase as more user-visible files are created. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .
files_private_used	integer	Number of system metadata files used. If the referenced file system is restricted or offline, a value of 0 is returned. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .

Name	Type	Description
files_total	integer	Maximum number of user-visible files that this referenced file system can currently hold. If the referenced file system is restricted or offline, a value of 0 is returned.
files_used	integer	Number of user-visible files used in the referenced file system. If the referenced file system is restricted or offline, a value of 0 is returned.
max_files_available	integer	The count of the maximum number of user-visible files currently allowable on the referenced file system.
max_files_possible	integer	The largest value to which the maxfiles-available parameter can be increased by reconfiguration, on the referenced file system.
max_files_used	integer	The number of user-visible files currently in use on the referenced file system.
used_percent	integer	The percentage of disk space currently in use based on user-visible file count on the referenced file system.
version	integer	The inofile-version of the aggregate. If the referenced file system is restricted or offline, a value of 0 is returned. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .

iops

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

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

latency

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

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

throughput

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

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

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

metric

The most recent sample of I/O metrics for the aggregate.

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

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.

node

Node where the aggregate currently resides.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

raid_group

Name	Type	Description
default	integer	Default number of disks in a RAID group.

Name	Type	Description
max	integer	Maximum number of disks allowed in a RAID group.
min	integer	Minimum number of disks allowed in a RAID group.

layout_requirement

Name	Type	Description
aggregate_min_disks	integer	Minimum number of disks to create an aggregate.
default	boolean	Indicates if this RAID type is the default.
raid_group	raid_group	
raid_type	string	RAID type.

node

Node where the spares are assigned.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

aggregate_spare

Name	Type	Description
checksum_style	string	The checksum type that has been assigned to the spares.
disk_class	string	Disk class of spares.
disk_type	string	Type of disk.
is_partition	boolean	Indicates whether a disk is partitioned (true) or whole (false)
layout_requirements	array[layout_requirement]	Available RAID protections and their restrictions.

Name	Type	Description
node	node	Node where the spares are assigned.
size	integer	Usable size of each spare, in bytes.
syncmirror_pool	string	SyncMirror spare pool.
total	integer	Total number of spares in the bucket. The total spare count for each class of spares also includes reserved spare capacity recommended by ONTAP best practices. <ul style="list-style-type: none"> • example: 10 • readOnly: 1 • Introduced in: 9.11 • x-nullable: true
usable	integer	Total number of usable spares in the bucket. The usable count for each class of spares does not include reserved spare capacity recommended by ONTAP best practices. <ul style="list-style-type: none"> • example: 9 • readOnly: 1 • Introduced in: 9.6 • x-nullable: true

snapshot

Name	Type	Description
files_total	integer	Total files allowed in Snapshot copies
files_used	integer	Total files created in Snapshot copies
max_files_available	integer	Maximum files available for Snapshot copies

Name	Type	Description
max_files_used	integer	Files in use by Snapshot copies

block_storage

Name	Type	Description
aggregate_metadata	integer	Space used by different metafiles and internal operations inside the aggregate, in bytes.
aggregate_metadata_percent	integer	Aggregate metadata as a percentage.
available	integer	Space available in bytes.
data_compacted_count	integer	Amount of compacted data in bytes.
data_compaction_space_saved	integer	Space saved in bytes by compacting the data.
data_compaction_space_saved_percent	integer	Percentage saved by compacting the data.
full_threshold_percent	integer	The aggregate used percentage at which 'monitor.volume.full' EMS is generated.
inactive_user_data	integer	The size that is physically used in the block storage and has a cold temperature, in bytes. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>block_storage.inactive_user_data</code> or <code>**</code> .

Name	Type	Description
inactive_user_data_percent	integer	The percentage of inactive user data in the block storage. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>block_storage.inactive_user_data_percent</code> or <code>**</code> .
physical_used	integer	Total physical used size of an aggregate in bytes.
physical_used_percent	integer	Physical used percentage.
size	integer	Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
used	integer	Space used or reserved in bytes. Includes volume guarantees and aggregate metadata.
used_including_snapshot_reserve	integer	Total used including the Snapshot copy reserve, in bytes.
used_including_snapshot_reserve_percent	integer	Total used including the Snapshot reserve as a percentage.
used_percent	integer	Aggregate used percentage.
volume_deduplication_shared_count	integer	Amount of shared bytes counted by storage efficiency.
volume_deduplication_space_saved	integer	Amount of space saved in bytes by storage efficiency.
volume_deduplication_space_saved_percent	integer	Percentage of space saved by storage efficiency.

Name	Type	Description
volume_footprints_percent	integer	A summation of volume footprints inside the aggregate, as a percentage. A volume's footprint is the amount of space being used for the volume in the aggregate.

cloud_storage

Name	Type	Description
used	integer	Used space in bytes in the cloud store. Only applicable for aggregates with a cloud store tier.

efficiency

Storage efficiency.

Name	Type	Description
auto_adaptive_compression_savings	boolean	Indicates whether or not aggregate has auto adaptive compression savings.
cross_volume_background_dedupe	boolean	Indicates whether or not cross volume background deduplication is enabled.
cross_volume_dedupe_savings	boolean	Indicates whether or not aggregate has cross volume deduplication savings.
cross_volume_inline_dedupe	boolean	Indicates whether or not cross volume inline deduplication is enabled.
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

efficiency_without_snapshots

Storage efficiency that does not include the savings provided by Snapshot copies.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

efficiency_without_snapshots_flexclones

Storage efficiency that does not include the savings provided by Snapshot copies and Flexclone volumes.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

snapshot

Name	Type	Description
available	integer	Available space for Snapshot copies in bytes
reserve_percent	integer	Percentage of space reserved for Snapshot copies
total	integer	Total space for Snapshot copies in bytes
used	integer	Space used by Snapshot copies in bytes
used_percent	integer	Percentage of disk space used by Snapshot copies

space

Name	Type	Description
block_storage	block_storage	

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

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.

Name	Type	Description
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.

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

The real time I/O statistics for the aggregate.

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.
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.

aggregate

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>_tags</code>	<code>array[string]</code>	Tags are an optional way to track the uses of a resource. Tag values must be formatted as <code>key:value</code> strings.
<code>block_storage</code>	<code>block_storage</code>	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
<code>cloud_storage</code>	<code>cloud_storage</code>	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
<code>create_time</code>	<code>string</code>	Timestamp of aggregate creation.
<code>data_encryption</code>	<code>data_encryption</code>	
<code>dr_home_node</code>	<code>dr_home_node</code>	Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.
<code>home_node</code>	<code>home_node</code>	Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.
<code>inactive_data_reporting</code>	<code>inactive_data_reporting</code>	
<code>inode_attributes</code>	<code>inode_attributes</code>	
<code>is_spare_low</code>	<code>boolean</code>	Specifies whether the aggregate is in a spares low condition on any of the RAID groups. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>footprint</code> or <code>**</code> .

Name	Type	Description
metric	metric	The most recent sample of I/O metrics for the aggregate.
name	string	Aggregate name.
node	node	Node where the aggregate currently resides.
recommendation_spares	array[aggregate_spare]	Information on the aggregate's remaining hot spare disks.
sidl_enabled	boolean	Specifies whether or not SIDL is enabled on the aggregate.
snaplock_type	string	SnapLock type.
snapshot	snapshot	
space	space	
state	string	Operational state of the aggregate.
statistics	statistics	The real time I/O statistics for the aggregate.
uuid	string	Aggregate UUID.
volume-count	integer	Number of volumes in the aggregate.

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.

action

Name	Type	Description
arguments	array[string]	Arguments present in the specified action message.

Name	Type	Description
code	integer	Corrective action code of the specified action.
message	string	Specifies the corrective action to be taken to resolve the issue.

warning

Name	Type	Description
arguments	array[string]	Arguments present in the warning message encountered.
code	integer	Warning code of the warning encountered.
message	string	Details of the warning encountered by the aggregate simulate query.

aggregate_warning

Name	Type	Description
action	action	
name	string	Name of the entity that returns the warning.
warning	warning	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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

Retrieve a collection of cloud stores used by an aggregate

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

Introduced In: 9.6

Retrieves the collection of cloud stores used by an aggregate.

Related ONTAP commands

- `storage aggregate object-store show`

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
availability	string	query	False	Filter by availability
primary	boolean	query	False	Filter by primary
unavailable_reason. message	string	query	False	Filter by unavailable_reason. message • Introduced in: 9.7
aggregate.name	string	query	False	Filter by aggregate.name • Introduced in: 9.9
used	integer	query	False	Filter by used
mirror_degraded	boolean	query	False	Filter by mirror_degraded

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

Response

Status: 200, Ok

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "aggregate": {
      "name": "aggr1"
    },
    "availability": "available",
    "target": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "target1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "unavailable_reason": {
      "message": "string"
    },
    "unreclaimed_space_threshold": 20,
    "used": 0
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

aggregate

Aggregate

Name	Type	Description
name	string	

target

Cloud target

Name	Type	Description
_links	_links	
name	string	
uuid	string	

unavailable_reason

Name	Type	Description
message	string	Indicates why the object store is unavailable.

cloud_store

Name	Type	Description
_links	_links	

Name	Type	Description
aggregate	aggregate	Aggregate
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a mirrored FabricPool.
target	target	Cloud target
unavailable_reason	unavailable_reason	
unreclaimed_space_threshold	integer	Usage threshold for reclaiming unused space in the cloud store. Valid values are 0 to 99. The default value depends on the provider type. This can be specified in PATCH but not POST.
used	integer	The amount of object space used. Calculated every 5 minutes and cached.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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

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

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

Introduced In: 9.6

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

Required properties

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

Recommended optional properties

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

Default property values

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

Related ONTAP commands

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

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID

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

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

Request Body

Name	Type	Description
_links	_links	
aggregate	aggregate	Aggregate
availability	string	Availability of the object store.

Name	Type	Description
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a mirrored FabricPool.
target	target	Cloud target
unavailable_reason	unavailable_reason	
unreclaimed_space_threshold	integer	Usage threshold for reclaiming unused space in the cloud store. Valid values are 0 to 99. The default value depends on the provider type. This can be specified in PATCH but not POST.
used	integer	The amount of object space used. Calculated every 5 minutes and cached.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregate": {
    "name": "aggr1"
  },
  "availability": "available",
  "target": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "target1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "unavailable_reason": {
    "message": "string"
  },
  "unreclaimed_space_threshold": 20,
  "used": 0
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Headers

Name	Description	Type
Location	Useful for tracking the resource location	string

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregate

Aggregate

Name	Type	Description
name	string	

target

Cloud target

Name	Type	Description
_links	_links	
name	string	
uuid	string	

unavailable_reason

Name	Type	Description
message	string	Indicates why the object store is unavailable.

cloud_store

Name	Type	Description
_links	_links	
aggregate	aggregate	Aggregate
availability	string	Availability of the object store.

Name	Type	Description
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a mirrored FabricPool.
target	target	Cloud target
unavailable_reason	unavailable_reason	
unreclaimed_space_threshold	integer	Usage threshold for reclaiming unused space in the cloud store. Valid values are 0 to 99. The default value depends on the provider type. This can be specified in PATCH but not POST.
used	integer	The amount of object space used. Calculated every 5 minutes and cached.

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

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

Remove a cloud target from an aggregate

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

Introduced In: 9.6

Removes the specified cloud target from the aggregate. Only removal of a mirror is allowed. The primary cannot be removed. This request starts a job and returns a link to that job.

Related ONTAP commands

- `storage aggregate object-store unmirror`

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
target.uuid	string	path	True	Cloud target UUID

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

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve the cloud store for an aggregate

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

Introduced In: 9.6

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

Related ONTAP commands

- `storage aggregate object-store show`

Parameters

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

Response

Status: 200, Ok

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

Name	Type	Description
unreclaimed_space_threshold	integer	Usage threshold for reclaiming unused space in the cloud store. Valid values are 0 to 99. The default value depends on the provider type. This can be specified in PATCH but not POST.
used	integer	The amount of object space used. Calculated every 5 minutes and cached.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregate": {
    "name": "aggr1"
  },
  "availability": "available",
  "target": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "target1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "unavailable_reason": {
    "message": "string"
  },
  "unreclaimed_space_threshold": 20,
  "used": 0
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregate

Aggregate

Name	Type	Description
name	string	

target

Cloud target

Name	Type	Description
_links	_links	
name	string	
uuid	string	

unavailable_reason

Name	Type	Description
message	string	Indicates why the object store is unavailable.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Update a cloud store

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

Introduced In: 9.6

Updates the cloud store specified by the UUID with the fields in the body. This request starts a job and returns a link to that job.

Related ONTAP commands

- `storage aggregate object-store modify`

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
target.uuid	string	path	True	Cloud target UUID

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

Request Body

Name	Type	Description
_links	_links	
aggregate	aggregate	Aggregate
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.

Name	Type	Description
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a mirrored FabricPool.
target	target	Cloud target
unavailable_reason	unavailable_reason	
unreclaimed_space_threshold	integer	Usage threshold for reclaiming unused space in the cloud store. Valid values are 0 to 99. The default value depends on the provider type. This can be specified in PATCH but not POST.
used	integer	The amount of object space used. Calculated every 5 minutes and cached.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregate": {
    "name": "aggr1"
  },
  "availability": "available",
  "target": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "target1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "unavailable_reason": {
    "message": "string"
  },
  "unreclaimed_space_threshold": 20,
  "used": 0
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregate

Aggregate

Name	Type	Description
name	string	

target

Cloud target

Name	Type	Description
_links	_links	
name	string	
uuid	string	

unavailable_reason

Name	Type	Description
message	string	Indicates why the object store is unavailable.

cloud_store

Name	Type	Description
_links	_links	
aggregate	aggregate	Aggregate
availability	string	Availability of the object store.

Name	Type	Description
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a mirrored FabricPool.
target	target	Cloud target
unavailable_reason	unavailable_reason	
unreclaimed_space_threshold	integer	Usage threshold for reclaiming unused space in the cloud store. Valid values are 0 to 99. The default value depends on the provider type. This can be specified in PATCH but not POST.
used	integer	The amount of object space used. Calculated every 5 minutes and cached.

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

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

Manage storage aggregate plexes

Storage aggregates aggregate.uuid plexes endpoint overview

Overview

The Storage Aggregate Plex API provides relevant state information for each plex in the aggregate. For each plex, details are provided for the RAID groups in the plex and the disks that make up each RAID group.

Examples

Retrieving all aggregates and plexes

The following example shows the response with a list of aggregates and plexes:

The call:

```
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/*/plexes" -H "accept: application/json"</mgmt-ip>
```

The response:

```
{ "records": [ { "aggregate": { "uuid": "04b7296e-a302-42a6-a2a9-dda6be054b29", "name": "test2" }, "name": "plex0" }, { "aggregate": { "uuid": "04b7296e-a302-42a6-a2a9-dda6be054b29", "name": "test2" }, "name": "plex1" }, { "aggregate": { "uuid": "66c4b221-65ff-4211-9b58-ada3c6fc41af", "name": "test" }, "name": "plex0" }, { "aggregate": { "uuid": "66c4b221-65ff-4211-9b58-ada3c6fc41af", "name": "test" }, "name": "plex1" }, { "aggregate": { "uuid": "7ee89e48-5d81-4609-9e1b-5d8d0995a886", "name": "aggr1" }, "name": "plex0" }, { "aggregate": { "uuid": "8bb2e3bf-c4f1-4748-9033-ca9231cf1c40", "name": "test3" }, "name": "plex0" }, { "aggregate": { "uuid": "8bb2e3bf-c4f1-4748-9033-ca9231cf1c40", "name": "test3" }, "name": "plex1" }, { "aggregate": { "uuid": "8f13de5c-99cf-4ada-884c-3cc32deb304a", "name": "aggr2" }, "name": "plex0" } ], "num_records": 8 }
```

```
### Retrieving a specific plex in all aggregates
```

The following example shows the response with a list of specific plexes in all aggregates:

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/*/plexes/plex0" -H "accept: application/json"
```

```

# The response:
{
  "records": [
    {
      "aggregate": {
        "uuid": "04b7296e-a302-42a6-a2a9-dda6be054b29",
        "name": "test2"
      },
      "name": "plex0"
    },
    {
      "aggregate": {
        "uuid": "66c4b221-65ff-4211-9b58-ada3c6fc41af",
        "name": "test"
      },
      "name": "plex0"
    },
    {
      "aggregate": {
        "uuid": "7ee89e48-5d81-4609-9e1b-5d8d0995a886",
        "name": "aggr1",
      },
      "name": "plex0",
    },
    {
      "aggregate": {
        "uuid": "8bb2e3bf-c4f1-4748-9033-ca9231cf1c40",
        "name": "test3",
      },
      "name": "plex0",
    },
    {
      "aggregate": {
        "uuid": "8f13de5c-99cf-4ada-884c-3cc32deb304a",
        "name": "aggr2",
      },
      "name": "plex0"
    }
  ],
  "num_records": 8
}

```

Retrieving the list of plexes in an aggregate

The following example shows the response with the list of plexes in an aggregate:

```
# The API:
/api/storage/aggregates/{uuid}/plexes

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-8f01-712f626c983c/plexes" -H "accept: application/json"

# The response:
{
  "records": [
    {
      "name": "plex0",
    },
    {
      "name": "plex4",
    }
  ],
  "num_records": 2,
}
```

Retrieving a specific plex in an aggregate

The following example shows the response when requesting a specific plex of an aggregate:

```
# The API:
/api/storage/aggregates/{uuid}/plexes/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-8f01-712f626c983c/plexes/plex0" -H "accept: application/json"

# The response:
{
  "aggregate": {
    "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
    "name": "test1",
  },
  "name": "plex0",
  "online": true,
  "state": "normal",
  "pool": "pool0",
  "resync": {
    "active": false
  },
  "raid_groups": [
```

```

{
  "name": "rg0",
  "raid_type": "raid_dp",
  "cache_tier": false,
  "degraded": false,
  "recomputing_parity": {
    "active": false
  },
  "reconstruct": {
    "active": false
  },
  "disks": [
    {
      "position": "dparity",
      "state": "normal",
      "type": "ssd",
      "usable_size": 86769664,
      "disk": {
        "name": "1.1.29",
      }
    },
    {
      "position": "parity",
      "state": "normal",
      "type": "ssd",
      "usable_size": 86769664,
      "disk": {
        "name": "1.1.4",
      }
    },
    {
      "position": "data",
      "state": "normal",
      "type": "ssd",
      "usable_size": 86769664,
      "disk": {
        "name": "1.1.30",
      }
    },
    {
      "position": "data",
      "state": "normal",
      "type": "ssd",
      "usable_size": 86769664,
      "disk": {
        "name": "1.1.5",
      }
    }
  ]
}

```



```

    }
  },
  {
    "position": "data",
    "state": "normal",
    "type": "ssd",
    "usable_size": 86769664,
    "disk": {
      "name": "1.1.31",
    }
  },
  {
    "position": "data",
    "state": "normal",
    "type": "ssd",
    "usable_size": 86769664,
    "disk": {
      "name": "1.1.6",
    }
  }
]
}
],
}

```

Retrieve a collection of plexes for an aggregate

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

Introduced In: 9.6

Retrieves the collection of plexes for the specified aggregate.

Related ONTAP commands

- storage aggregate plex show

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
resync.percent	integer	query	False	Filter by resync.percent

Name	Type	In	Required	Description
resync.active	boolean	query	False	Filter by resync.active
resync.level	string	query	False	Filter by resync.level
raid_groups.cache_tier	boolean	query	False	Filter by raid_groups.cache_tier • Introduced in: 9.7
raid_groups.reconstruct.active	boolean	query	False	Filter by raid_groups.reconstruct.active • Introduced in: 9.7
raid_groups.reconstruct.percent	integer	query	False	Filter by raid_groups.reconstruct.percent • Introduced in: 9.7
raid_groups.disks.state	string	query	False	Filter by raid_groups.disks.state • Introduced in: 9.7
raid_groups.disks.position	string	query	False	Filter by raid_groups.disks.position • Introduced in: 9.7
raid_groups.disks.type	string	query	False	Filter by raid_groups.disks.type • Introduced in: 9.7

Name	Type	In	Required	Description
raid_groups.disks.usable_size	integer	query	False	Filter by raid_groups.disks.usable_size • Introduced in: 9.7
raid_groups.disks.disk.name	string	query	False	Filter by raid_groups.disks.disk.name • Introduced in: 9.7
raid_groups.recomputing_parity.active	boolean	query	False	Filter by raid_groups.recomputing_parity.active • Introduced in: 9.7
raid_groups.recomputing_parity.percent	integer	query	False	Filter by raid_groups.recomputing_parity.percent • Introduced in: 9.7
raid_groups.name	string	query	False	Filter by raid_groups.name • Introduced in: 9.7
raid_groups.degraded	boolean	query	False	Filter by raid_groups.degraded • Introduced in: 9.7
raid_groups.raid_type	string	query	False	Filter by raid_groups.raid_type • Introduced in: 9.9

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

Response

Status: 200, Ok

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "num_records": 1,
  "records": {
    "aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "name": "plex0",
    "pool": "pool0",
    "raid_groups": {
      "disks": {
        "disk": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "1.0.1"
        }
      },
      "position": "data",
      "state": "normal",
      "type": "ssd",

```

```

    "usable_size": 947912704
  },
  "name": "rg0",
  "raid_type": "raid_dp",
  "recomputing_parity": {
    "percent": 10
  },
  "reconstruct": {
    "percent": 10
  }
},
"resync": {
  "level": "full",
  "percent": 10
},
"state": "normal"
}
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

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

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

_links

Name	Type	Description
self	href	

aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

disk

Name	Type	Description
_links	_links	
name	string	

raid_group_disk

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

recomputing_parity

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

reconstruct

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

raid_group

Name	Type	Description
cache_tier	boolean	RAID group is a cache tier

Name	Type	Description
degraded	boolean	RAID group is degraded. A RAID group is degraded when at least one disk from that group has failed or is offline.
disks	array[raid_group_disk]	
name	string	RAID group name
raid_type	string	RAID type of the raid group.
recomputing_parity	recomputing_parity	
reconstruct	reconstruct	

resync

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

plex

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

Retrieve a plex specified by the aggregate UUID and plex name

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

Introduced In: 9.6

Retrieves the plex specified by the aggregate UUID and plex name.

Related ONTAP commands

- `storage aggregate plex show`

Parameters

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

Response

Status: 200, Ok

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

Example response

```
{
  "aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "plex0",
  "pool": "pool0",
  "raid_groups": {
    "disks": {
      "disk": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "1.0.1"
      },
      "position": "data",
      "state": "normal",
      "type": "ssd",
      "usable_size": 947912704
    },
    "name": "rg0",
    "raid_type": "raid_dp",
    "recomputing_parity": {
      "percent": 10
    },
    "reconstruct": {
      "percent": 10
    }
  },
  "resync": {
    "level": "full",
    "percent": 10
  },
  "state": "normal"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

disk

Name	Type	Description
_links	_links	
name	string	

raid_group_disk

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

recomputing_parity

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

reconstruct

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

raid_group

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

resync

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage storage aggregates

Storage aggregates UUID endpoint overview

Updating storage aggregates

The PATCH operation is used to modify properties of the aggregate. There are several properties that can be modified on an aggregate. Only one property can be modified for each PATCH request. PATCH operations on the aggregate's disk count will be blocked while one or more nodes in the cluster are simulating or implementing automatic aggregate creation.

The following is a list of properties that can be modified using the PATCH operation including a brief description for each:

- name - This property can be changed to rename the aggregate.
- node.name and node.uuid - Either property can be updated in order to relocate the aggregate to a different node in the cluster.
- state - This property can be changed to 'online' or 'offline'. Setting an aggregate 'offline' would automatically offline all the volumes currently hosted on the aggregate.
- block_storage.mirror.enabled - This property can be changed from 'false' to 'true' in order to mirror the aggregate, if the system is capable of doing so.
- block_storage.primary.disk_count - This property can be updated to increase the number of disks in an aggregate.
- block_storage.primary.raid_size - This property can be updated to set the desired RAID size.
- block_storage.primary.raid_type - This property can be updated to set the desired RAID type.
- cloud_storage.tiering_fullness_threshold - This property can be updated to set the desired tiering fullness threshold if using FabricPool.

- `cloud_storage.migrate_threshold` - This property can be updated to set the desired migrate threshold if using FabricPool.
- `data_encryption.software_encryption_enabled` - This property enables or disables NAE on the aggregate.
- `block_storage.hybrid_cache.storage_pools.allocation_units_count` - This property can be updated to add a storage pool to the aggregate specifying the number of allocation units.
- `block_storage.hybrid_cache.storage_pools.name` - This property can be updated to add a storage pool to the aggregate specifying the storage pool name. `block_storage.hybrid_cache.storage_pools.uuid` or this field must be specified with `block_storage.hybrid_cache.storage_pools.allocation_units_count`.
- `block_storage.hybrid_cache.storage_pools.uuid` - This property can be updated to add a storage pool to the aggregate specifying the storage pool uuid. `block_storage.hybrid_cache.storage_pools.name` or this field must be specified with `block_storage.hybrid_cache.storage_pools.allocation_units_count`.
- `block_storage.hybrid_cache.raid_size` - This property can be updated to set the desired RAID size. This property can also be specified on the first time addition of a storage pool to the aggregate.
- `block_storage.hybrid_cache.raid_type` - This property can be updated to set the desired RAID type of a physical SSD Flash Pool. This property can also be specified on the first time addition of a storage pool to the aggregate. When specifying a raidtype of raid4, the node is required to have spare SSDs for the storage pool as well.
- `block_storage.hybrid_cache.disk_count` - This property can be specified on the first time addition of physical SSD cache to the aggregate. It can also be updated to increase the number of disks in the physical SSD cache of a hybrid aggregate.

Aggregate expansion

The PATCH operation also supports automatically expanding an aggregate based on the spare disks which are present within the system. Running PATCH with the query "auto_provision_policy" set to "expand" starts the recommended expansion job. In order to see the expected change in capacity before starting the job, call GET on an aggregate instance with the query "auto_provision_policy" set to "expand".

Manual simulated aggregate expansion

The PATCH operation also supports simulated manual expansion of an aggregate. Running PATCH with the query "simulate" set to "true" and "block_storage.primary.disk_count" set to the final disk count will start running the prechecks associated with expanding the aggregate to the proposed size. The response body will include information on how many disks the aggregate can be expanded to, any associated warnings, along with the proposed final size of the aggregate.

Deleting storage aggregates

If volumes exist on an aggregate, they must be deleted or moved before the aggregate can be deleted. See the /storage/volumes API for details on moving or deleting volumes.

Adding a storage pool to an aggregate

A storage pool can be added to an aggregate by patching the field "block_storage.hybrid_cache.storage_pools.allocation_units_count" while also specifying the specific storage pool using the "block_storage.hybrid_cache.storage_pools.name" or "block_storage.hybrid_cache.storage_pools.uuid". Subsequent patches to the aggregate can be completed to increase allocation unit counts or adding additional storage pools. On the first time addition of a storage pool to the aggregate, the raidtype can be optionally specified using the "block_storage.hybrid_cache.raid_type" field.

Adding physical SSD cache capacity to an aggregate

The PATCH operation supports addition of a new physical SSD cache to an aggregate. It also supports expansion of existing physical SSD cache in the hybrid aggregate. Running PATCH with "block_storage.hybrid_cache.disk_count" set to the final disk count will expand the physical SSD cache of the hybrid aggregate to the proposed size. The RAID type can be optionally specified using the "block_storage.hybrid_cache.raid_type" field. The RAID size can be optionally specified using the "block_storage.hybrid_cache.raid_size" field. These operations can also be simulated by setting the query "simulate" to "true".

Examples

Retrieving a specific aggregate from the cluster

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

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

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/870dd9f2-bdfa-4167-
b692-57d1cec874d4" -H "accept: application/json"

# The response:
{
  "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
  "name": "test1",
  "node": {
    "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
    "name": "node-1",
  },
  "home_node": {
    "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
    "name": "node-1",
  },
  "space": {
    "block_storage": {
      "size": 235003904,
      "available": 191942656,
      "used": 43061248,
      "full_threshold_percent": 98,
      "physical_used": 5271552,
      "physical_used_percent": 1,
      "volume_footprints_percent": 14,
      "aggregate_metadata": 2655,
      "aggregate_metadata_percent": 8,
```

```

    "used_including_snapshot_reserve": 674685,
    "used_including_snapshot_reserve_percent": 35,
    "data_compacted_count": 666666,
    "data_compaction_space_saved": 654566,
    "data_compaction_space_saved_percent": 47,
    "volume_deduplication_shared_count": 567543,
    "volume_deduplication_space_saved": 23765,
    "volume_deduplication_space_saved_percent": 32,
    "used_percent": 50,
  },
  "snapshot": {
    "used_percent": 45,
    "available": 2000,
    "total": 5000,
    "used": 3000,
    "reserve_percent": 20
  },
  "cloud_storage": {
    "used": 0
  },
  "efficiency": {
    "savings": 1408029,
    "ratio": 6.908119720880661,
    "logical_used": 1646350,
    "cross_volume_background_dedupe": true,
    "cross_volume_inline_dedupe": false,
    "cross_volume_dedupe_savings": true,
    "auto_adaptive_compression_savings": false
  },
  "efficiency_without_snapshots": {
    "savings": 0,
    "ratio": 1,
    "logical_used": 737280
  },
  "efficiency_without_snapshots_flexclones": {
    "savings": 5000,
    "ratio": 2,
    "logical_used": 10000
  }
},
"snapshot": {
  "files_total": 10,
  "files_used": 3,
  "max_files_available": 5,
  "max_files_used": 50
},

```

```
"state": "online",
"snaplock_type": "non_snaplock",
"create_time": "2018-12-04T15:40:38-05:00",
"data_encryption": {
  "software_encryption_enabled": false,
  "drive_protection_enabled": false
},
"block_storage": {
  "uses_partitions": false,
  "storage_type": "vmdisk",
  "primary": {
    "disk_count": 6,
    "disk_class": "solid_state",
    "raid_type": "raid_dp",
    "raid_size": 24,
    "checksum_style": "block",
    "disk_type": "ssd"
  },
  "hybrid_cache": {
    "enabled": false
  },
  "mirror": {
    "enabled": false,
    "state": "unmirrored"
  },
  "plexes": [
    {
      "name": "plex0",
    }
  ]
},
"cloud_storage": {
  "attach_eligible": false
},
"inode_attributes": {
  "files_total": 31136,
  "files_used": 97,
  "max_files_available": 31136,
  "max_files_possible": 2844525,
  "max_files_used": 97,
  "used_percent": 5
},
"volume_count": 0,
}
```

Retrieving statistics and metric for an aggregate

In this example, the API returns the "statistics" and "metric" properties for the aggregate requested.

```
#The API:
/api/storage/aggregates/{uuid}?fields=statistics,metric

#The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/538bf337-1b2c-11e8-
bad0-005056b48388?fields=statistics,metric" -H "accept: application/json"

#The response:
{
  "uuid": "538bf337-1b2c-11e8-bad0-005056b48388",
  "name": "aggr4",
  "metric": {
    "timestamp": "2019-07-08T22:16:45Z",
    "duration": "PT15S",
    "status": "ok",
    "throughput": {
      "read": 7099,
      "write": 840226,
      "other": 193293789,
      "total": 194141115
    },
    "latency": {
      "read": 149,
      "write": 230,
      "other": 123,
      "total": 124
    },
    "iops": {
      "read": 1,
      "write": 17,
      "other": 11663,
      "total": 11682
    },
  },
  "statistics": {
    "timestamp": "2019-07-08T22:17:09Z",
    "status": "ok",
    "throughput_raw": {
      "read": 3106045952,
      "write": 63771742208,
      "other": 146185560064,
      "total": 213063348224
    }
  }
}
```

```

    },
    "latency_raw": {
      "read": 54072313,
      "write": 313354426,
      "other": 477201985,
      "total": 844628724
    },
    "iops_raw": {
      "read": 328267,
      "write": 1137230,
      "other": 1586535,
      "total": 3052032
    }
  },
}

```

For more information and examples on viewing historical performance metrics for any given aggregate, see [DOC /storage/aggregates/{uuid}/metrics](#)

Simulating aggregate expansion

The following example shows the response for a simulated data aggregate expansion based on the values of the 'block_storage.primary.disk_count' attribute passed in. The query does not modify the existing aggregate but returns how the aggregate will look after the expansion along with any associated warnings. Simulated data aggregate expansion will be blocked while one or more nodes in the cluster are simulating or implementing automatic aggregate creation. This will be reflected in the following attributes:

- space.block_storage.size - Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
- block_storage.primary.disk_count - Number of disks that could be used to create the aggregate.

```

# The API:
/api/storage/aggregates/{uuid}?simulate=true

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/cae60cfe-deae-42bd-babb-ef437d118314?simulate=true" -H "accept: application/json" -d '{"block_storage": {"primary": {"disk_count": 13}}}'

# The response:
{
  "warnings": [
    {
      "name": "node_2_SSD_1",
      "warning": {
        "message": "One or more disks will not be added. 10 disks specified, 9 disks will be added."
      }
    }
  ]
}

```

```

    "code": 787170,
    "arguments": [
      "10",
      "9"
    ]
  }
},
"num_records": 1,
"records": [
  {
    "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
    "name": "node_2_SSD_1",
    "node": {
      "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
      "name": "node-2"
    },
    "space": {
      "block_storage": {
        "size": 1116180480
      }
    },
    "block_storage": {
      "primary": {
        "disk_count": 12,
        "disk_class": "solid_state",
        "raid_type": "raid_dp",
        "disk_type": "ssd",
        "raid_size": 12,
        "simulated_raid_groups": [
          {
            "name": "test/plex0/rg0",
            "existing_parity_disk_count": 2,
            "added_parity_disk_count": 0,
            "existing_data_disk_count": 1,
            "added_data_disk_count": 9,
            "usable_size": 12309487,
            "is_partition": false
          }
        ]
      },
      "hybrid_cache": {
        "enabled": false
      },
      "mirror": {
        "enabled": false
      }
    }
  }
]

```



```

    }
  },
}
]
}

```

Manual aggregate expansion with disk size query

The following example shows the response for aggregate expansion based on the values of the 'block_storage.hybrid_cache.disk_count' attribute based on the disk size passed in.

```

# The API:
/api/storage/aggregate/{uuid}?disk_size={disk_size}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/cae60cfe-deae-42bd-babb-ef437d118314?disk_size=1902379008" -H "accept: application/json" -d "{\"block_storage\": {\"hybrid_cache\": {\"disk_count\": 4}}}"

# The response:
{
  "job": {
    "uuid": "c103d15e-730b-11e8-a57f-005056b465d6",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/c103d15e-730b-11e8-a57f-005056b465d6"
      }
    }
  }
}

```

Simulating a manual aggregate expansion with disk size query

The following example shows the response for a manual aggregate expansion based on the values of the 'block_storage.hybrid_cache.disk_count' attribute based on the disk size passed in. The query internally maps out the appropriate expansion as well as warnings that may be associated for the hybrid enabled aggregate.

```

# The API:
/api/storage/aggregate/{uuid}?simulate=true&disk_size=1902379008

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/cae60cfe-deae-42bd-babb-ef437d118314?simulate=true&disk_size=1902379008" -H "accept: application/json" -d "{\"block_storage\": {\"hybrid_cache\": {\"disk_count\": 4}}}"

```

```

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
      "name": "ag1",
      "node": {
        "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
        "name": "node-2",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/4046dda8-f802-11e8-8f6d-005056bb2030"
          }
        }
      },
      "block_storage": {
        "primary": {
          "disk_count": 4,
          "disk_class": "virtual",
          "raid_type": "raid_dp",
          "disk_type": "vm_disk",
        },
        "hybrid_cache": {
          "disk_type": "ssd",
          "enabled": true,
          "disk_count": 4,
          "raid_type": "raid_dp",
          "size": 3761766400,
          "simulated_raid_groups": [
            {
              "name": "test/plex0/rg0",
              "existing_parity_disk_count": 2,
              "existing_data_disk_count": 1,
              "added_parity_disk_count": 0,
              "added_data_disk_count": 1,
              "usable_size": 1880883200,
              "is_partition": false
            }
          ],
        },
        "mirror": {
          "enabled": false
        },
        "_links": {

```

```

        "self": {
            "href": "/api/storage/aggregates/cae60cfe-deae-42bd-babb-
ef437d118314"
        }
    }
}
]
}

```

Simulating a manual aggregate expansion with raid group query

The following example shows the response for a manual aggregate expansion based on the values of the 'block_storage.primary.disk_count' attribute passed in. The query internally maps out the appropriate expansion as well as warnings that may be associated and lays out the new raidgroups in a more detailed view. An additional query can be passed in to specify raidgroup addition by new raidgroup, all raidgroups or a specific raidgroup.

```

# The API:
/api/storage/aggregate/{uuid}?simulate=true&raid_group=[new&#124;all&#124;
rgX]

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/cae60cfe-deae-
42bd-babb-ef437d118314?simulate=true&raid_group=new" -H "accept:
application/json" -d "{\"block_storage\": {\"primary\": {\"disk_count\":
24}}}"

# The response:
{
  "warnings": [
    {
      "name": "test",
      "warning": {
        "code": 11,
        "message": "Number of unassigned disks attached to node \"node-2\":
6.",
        "arguments": [
          "6",
          "node-2"
        ]
      }
    }
  ],
  "num_records": 1,
  "records": [

```

```

{
  "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
  "name": "test",
  "node": {
    "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
    "name": "node-2"
  },
  "space": {
    "block_storage": {
      "size": 33292025856
    }
  },
  "block_storage": {
    "primary": {
      "disk_count": 24,
      "disk_class": "solid_state",
      "raid_type": "raid_dp",
      "disk_type": "ssd",
      "raid_size": 24,
      "simulated_raid_groups": [
        {
          "name": "test/plex0/rg0",
          "existing_parity_disk_count": 0,
          "added_parity_disk_count": 2,
          "existing_data_disk_count": 0,
          "added_data_disk_count": 10,
          "usable_size": 12309487,
          "is_partition": false
        },
        {
          "name": "test/plex1/rg1",
          "existing_parity_disk_count": 0,
          "added_parity_disk_count": 2,
          "existing_data_disk_count": 0,
          "added_data_disk_count": 10,
          "usable_size": 12309487,
          "is_partition": false
        }
      ]
    },
    "hybrid_cache": {
      "enabled": false
    },
    "mirror": {
      "enabled": false
    }
  }
}

```

```
    }  
  }  
]  
}
```

Retrieving the usable spare information for the cluster

The following example shows the response from retrieving usable spare information for the expansion of this particular aggregate. The output is restricted to only spares that are compatible with this aggregate.

```

# The API:
/api/storage/aggregates?show_spares=true&uuid={uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates?uuid=cae60cfe-deae-42bd-babb-ef437d118314&show_spares=true" -H "accept: application/json"

# The response:
{
  "records": [],
  "num_records": 0,
  "spares": [
    {
      "node": {
        "uuid": "0cdd84fa-b99c-11eb-b0ed-005056bb4fc2",
        "name": "node-2"
      },
      "disk_class": "solid_state",
      "disk_type": "ssd",
      "size": 3720609792,
      "checksum_style": "block",
      "syncmirror_pool": "pool0",
      "usable": 12,
      "layout_requirements": [
        {
          "raid_type": "raid_dp",
          "default": true,
          "aggregate_min_disks": 3,
          "raid_group": {
            "min": 3,
            "max": 28,
            "default": 24
          }
        }
      ]
    }
  ]
}

```

Retrieving the SSD spare count for the cluster

The following example shows the response from retrieving SSD spare count information for the expansion of this particular aggregate's hybrid cache tier. The output is restricted to only spares that are compatible with this aggregate.

```

# The API:
/api/storage/aggregates?show_spare=true&uuid={uuid}&flash_pool_eligible=true

# The response:
{
  "records": [],
  "num_records": 0,
  "spares": [
    {
      "node": {
        "uuid": "c35c5975-cbcb-11ec-a3e1-005056bbdb46",
        "name": "node-2"
      },
      "disk_class": "solid_state",
      "disk_type": "ssd",
      "size": 1902379008,
      "checksum_style": "block",
      "syncmirror_pool": "pool0",
      "is_partition": false,
      "usable": 1,
      "layout_requirements": [
        {
          "raid_type": "raid4",
          "default": true,
          "aggregate_min_disks": 2,
          "raid_group": {
            "min": 2,
            "max": 14,
            "default": 8
          }
        }
      ]
    }
  ]
}

```

Retrieving a recommendation for an aggregate expansion

The following example shows the response with the recommended data aggregate expansion based on what disks are present within the system. The query does not modify the existing aggregate but returns how the aggregate will look after the expansion. The recommendation will be reflected in the attributes - 'space.block_storage.size' and 'block_storage.primary.disk_count'. Recommended data aggregate expansion will be blocked while one or more nodes in the cluster are simulating or implementing automatic aggregate creation.

```

# The API:
/api/storage/aggregates/{uuid}?auto_provision_policy=expand

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/cae60cfe-deae-42bd-
babb-ef437d118314?auto_provision_policy=expand" -H "accept:
application/json"

# The response:
{
  "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
  "name": "node_2_SSD_1",
  "node": {
    "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
    "name": "node-2"
  },
  "space": {
    "block_storage": {
      "size": 1116180480
    }
  },
  "block_storage": {
    "primary": {
      "disk_count": 12,
      "disk_class": "solid_state",
      "raid_type": "raid_dp",
      "disk_type": "ssd",
      "raid_size": 24,
      "simulated_raid_groups": [
        {
          "name": "test/plex0/rg0",
          "parity_disk_count": 2,
          "data_disk_count": 10,
          "usable_size": 12309487,
          "is_partition": false
        }
      ]
    },
    "hybrid_cache": {
      "enabled": false
    },
    "mirror": {
      "enabled": false
    }
  }
}

```


Updating an aggregate in the cluster

The following example shows the workflow of adding disks to the aggregate.

Step 1: Check the current disk count on the aggregate.

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

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-8f01-712f626c983c?fields=block_storage.primary.disk_count" -H "accept: application/json"

# The response:
{
  "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
  "name": "test1",
  "block_storage": {
    "primary": {
      "disk_count": 6
    }
  }
},
}
```

Step 2: Update the aggregate with the new disk count in 'block_storage.primary.disk_count'. The response to PATCH is a job unless the request is invalid.

```

# The API:
/api/storage/aggregates

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-8f01-712f626c983c" -H "accept: application/hal+json" -d
"{\"block_storage\": {\"primary\": {\"disk_count\": 8}}}"

# The response:
{
  "job": {
    "uuid": "c103d15e-730b-11e8-a57f-005056b465d6",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/c103d15e-730b-11e8-a57f-005056b465d6"
      }
    }
  }
}

```

Step 3: Wait for the job to finish, then call GET to see the reflected change.

```

# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-8f01-712f626c983c?fields=block_storage.primary.disk_count" -H "accept: application/json"

# The response:
{
  "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
  "name": "test1",
  "block_storage": {
    "primary": {
      "disk_count": 8
    }
  },
}

```

Adding a storage pool to an aggregate

The following example shows how to add cache capacity from an existing storage pool to an aggregate. Step 1: Update the aggregate with the new storage pool allocation unit in

'block_storage.hybrid_cache.storage_pools.allocation_units_count'. Additionally, specify 'block_storage.hybrid_cache.storage_pools.name' or 'block_storage.hybrid_cache.storage_pools.uuid' to the storage pool. On the first storage pool, 'block_storage.hybrid_cache.raid_type' can be specified for the raidtype of the hybrid cache. The response to PATCH is a job unless the request is invalid.

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-8f01-712f626c983c" -H "accept: application/json" -d
"{\"block_storage\": {\"hybrid_cache\": {\"raid_type\": \"raid_dp\",
\"storage_pools\": [{ \"allocation_units_count\": 2, \"storage_pool\": {
\"name\": \"sp1\"}]}}}"

# The response:
{
  "job": {
    "uuid": "c103d15e-730b-11e8-a57f-005056b465d6",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/c103d15e-730b-11e8-a57f-005056b465d6"
      }
    }
  }
}
```

Step 2: Wait for the job to finish, then call GET to see the reflected change.

```

# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-8f01-712f626c983c?fields=block_storage.hybrid_cache" -H "accept: application/json"

# The response:
{
  "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
  "name": "test1",
  "hybrid_cache": {
    "enabled": true,
    "disk_count": 3,
    "raid_size": 24,
    "raid_type": "raid_dp",
    "size": 880279552,
    "used": 73728,
    "storage_pools": [
      {
        "allocation_units_count": 2,
        "storage_pool": {
          "name": "sp1",
          "uuid": "eeef0b24-846b-11ec-8fcb-005056bb12c7"
        }
      }
    ]
  }
}

```

Adding physical SSD cache capacity to an aggregate

The following example shows how to add physical SSD cache capacity to an aggregate. Step 1: Specify the number of disks to be added to cache in 'block_storage.hybrid_cache.disk_count'.

'block_storage.hybrid_cache.raid_type' can be specified for the RAID type of the hybrid cache.

'block_storage.hybrid_cache.raid_size' can be specified for the RAID size of the hybrid cache. The response to PATCH is a job unless the request is invalid.

```

# The API:
/api/storage/aggregates

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/caa8a9f1-0219-4eaf-bcad-e29c05042fe1" -H "accept: application/json" -d
'{"block_storage.hybrid_cache.disk_count":3,"block_storage.hybrid_cache.raid_type":"raid4"}'

# The response:
{
  "job": {
    "uuid": "c103d15e-730b-11e8-a57f-005056b465d6",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/c103d15e-730b-11e8-a57f-005056b465d6"
      }
    }
  }
}

```

Step 2: Wait for the job to finish, then call GET to see the reflected change.

```

# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/caa8a9f1-0219-4eaf-bcad-e29c05042fe1?fields=block_storage.hybrid_cache" -H "accept: application/json"

# The response:
{
  "uuid": "caa8a9f1-0219-4eaf-bcad-e29c05042fe1",
  "name": "test1",
  "hybrid_cache": {
    "enabled": true,
    "disk_count": 3,
    "raid_size": 24,
    "raid_type": "raid4",
    "size": 880279552,
    "used": 73728
  }
}

```

Simulated addition of physical SSD cache capacity to an aggregate

The following example shows the response for a simulated addition of physical SSD cache capacity to an aggregate based on the values of the 'block_storage.hybrid_cache.disk_count', 'block_storage.hybrid_cache.raid_type' and 'block_storage.hybrid_cache.raid_size' attributes passed in. The query does not modify the existing aggregate but returns how the aggregate will look after the expansion along with any associated warnings. Simulated addition of physical SSD cache capacity to an aggregate will be blocked while one or more nodes in the cluster are simulating or implementing automatic aggregate creation. This will be reflected in the following attributes:

- block_storage.hybrid_cache.size - Total usable cache space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
- block_storage.hybrid_cache.disk_count - Number of disks that can be added to the aggregate's cache tier.

```
# The API:
/api/storage/aggregates/{uuid}?simulate=true

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/7eb630d1-0e55-4cb6-8d90-957d6f4db54e?simulate=true" -H "accept: application/json" -d '{"block_storage.hybrid_cache.disk_count":6,"block_storage.hybrid_cache.raid_type":"raid4","block_storage.hybrid_cache.raid_size":3}'

# The response:
{
  "warnings": [
    {
      "name": "test",
      "warning": {
        "code": 18316,
        "message": "Operation will lead to creation of new raid group"
      }
    }
  ],
  "num_records": 1,
  "records": [
    {
      "uuid": "7eb630d1-0e55-4cb6-8d90-957d6f4db54e",
      "name": "test",
      "node": {
        "uuid": "30d69eb5-f046-11ec-9bba-005056bba492",
        "name": "node-1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/30d69eb5-f046-11ec-9bba-005056bba492"
          }
        }
      }
    }
  ]
}
```

```

},
"space": {
  "block_storage": {
    "size": 833777664
  }
},
"block_storage": {
  "primary": {
    "disk_count": 3,
    "disk_class": "virtual",
    "raid_type": "raid_dp",
    "disk_type": "vm_disk"
  },
  "hybrid_cache": {
    "disk_class": "solid_state",
    "disk_type": "ssd",
    "enabled": false,
    "disk_count": 6,
    "raid_type": "raid4",
    "size": 6771179520,
    "simulated_raid_groups": [
      {
        "name": "/test/plex0/rg1",
        "existing_parity_disk_count": 0,
        "existing_data_disk_count": 0,
        "added_parity_disk_count": 1,
        "added_data_disk_count": 2,
        "usable_size": 1880883200,
        "is_partition": false
      },
      {
        "name": "/test/plex0/rg2",
        "existing_parity_disk_count": 0,
        "existing_data_disk_count": 0,
        "added_parity_disk_count": 1,
        "added_data_disk_count": 2,
        "usable_size": 1880883200,
        "is_partition": false
      }
    ]
  },
  "mirror": {
    "enabled": false
  }
},
"_links": {

```

```
    "self": {
      "href": "/api/storage/aggregates/7eb630d1-0e55-4cb6-8d90-
957d6f4db54e"
    }
  }
}
]
```

The following example shows the workflow to enable software encryption on an aggregate.

Step 1: Check the current software encryption status of the aggregate.

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

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/f3aafdc6-be35-4d93-
9590-5a402bffb4b?fields=data_encryption.software_encryption_enabled" -H
"accept: application/json"

# The response:
{
  "uuid": "f3aafdc6-be35-4d93-9590-5a402bffb4b",
  "name": "aggr5",
  "data_encryption": {
    "software_encryption_enabled": false
  },
}
```

Step 2: Update the aggregate with the encryption status in 'data_encryption.software_encryption_enabled'. The response to PATCH is a job unless the request is invalid.


```

# The API:
/api/storage/aggregates

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/f3aafdc6-be35-4d93-9590-5a402bffbe4b" -H "accept: application/hal+json" -d '{"data_encryption": {"software_encryption_enabled": "true"}}'

# The response:
{
  "job": {
    "uuid": "6b7ab28e-168d-11ea-8a50-0050568eca76",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/6b7ab28e-168d-11ea-8a50-0050568eca76"
      }
    }
  }
}

```

Step 3: Wait for the job to finish, then call GET to see the reflected change.

```

# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/f3aafdc6-be35-4d93-9590-5a402bffbe4b?fields=data_encryption.software_encryption_enabled" -H "accept: application/json"

# The response:
{
  "uuid": "f3aafdc6-be35-4d93-9590-5a402bffbe4b",
  "name": "aggr5",
  "data_encryption": {
    "software_encryption_enabled": true
  },
}

```

Delete an aggregate specified by the UUID

```
DELETE /storage/aggregates/{uuid}
```

Introduced In: 9.6

Deletes the aggregate specified by the UUID. This request starts a job and returns a link to that job.

Related ONTAP commands

- `storage aggregate delete`

Parameters

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

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
460770	The aggregate delete job failed to delete the aggregate.
460777	Failed to get information on the delete job.
786435	Internal Error. Failed to create a communication handle.
786451	Failed to delete specified aggregate.
786468	VLDB is offline.
786472	Node that hosts the aggregate is offline.
786497	Cannot delete an aggregate that has volumes.
786771	Aggregate does not exist.
786867	Specified aggregate resides on the remote cluster.
786897	Specified aggregate cannot be deleted as it is a switched-over root aggregate.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve an aggregate specified by the UUID

GET /storage/aggregates/{uuid}

Introduced In: 9.6

Retrieves the aggregate specified by the UUID. The recommend query cannot be used for this operation.

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.

- `metric.*`
- `space.block_storage.inactive_user_data`
- `space.block_storage.inactive_user_data_percent`
- `space.footprint`
- `is_spare_low`
- `statistics.*`

Related ONTAP commands

- `storage aggregate show`

Parameters

Name	Type	In	Required	Description
<code>uuid</code>	string	path	True	Aggregate UUID
<code>show_spare</code>	boolean	query	False	If set to 'true', the spares object is returned instead of records to show the spare groups in the cluster that are compatible with the aggregate. The default setting is 'false'. • Introduced in: 9.12

Name	Type	In	Required	Description
flash_pool_eligible	boolean	query	False	If set to 'true' along with show_spares, the spares object is restricted to return spare groups that are compatible with flash pool creation or addition for the specified aggregate. The default setting is 'false'. <ul style="list-style-type: none"> Introduced in: 9.12
auto_provision_policy	string	query	False	If set to expand, a query is run on the system for the recommended optimal expansion layout of the aggregate. <ul style="list-style-type: none"> Introduced in: 9.8
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.
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.

Name	Type	Description
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation.
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	home_node	Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.
inactive_data_reporting	inactive_data_reporting	
inode_attributes	inode_attributes	
is_spare_low	boolean	Specifies whether the aggregate is in a spares low condition on any of the RAID groups. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .
metric	metric	The most recent sample of I/O metrics for the aggregate.
name	string	Aggregate name.
node	node	Node where the aggregate currently resides.
recommendation_spares	array[aggregate_spare]	Information on the aggregate's remaining hot spare disks.
sidl_enabled	boolean	Specifies whether or not SIDL is enabled on the aggregate.

Name	Type	Description
snaplock_type	string	SnapLock type.
snapshot	snapshot	
space	space	
state	string	Operational state of the aggregate.
statistics	statistics	The real time I/O statistics for the aggregate.
uuid	string	Aggregate UUID.
volume-count	integer	Number of volumes in the aggregate.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "disk_type": "fc",
      "raid_size": 24,
      "raid_type": "raid_dp",
      "simulated_raid_groups": {
      },
      "size": 1612709888,
      "storage_pools": {
        "storage_pool": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "storage_pool_1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "used": 26501122
    },
    "mirror": {
      "enabled": "",
      "state": "unmirrored"
    },
    "plexes": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "plex0"
    }
  },
}
```

```

"primary": {
  "checksum_style": "block",
  "disk_class": "performance",
  "disk_count": 8,
  "disk_type": "fc",
  "raid_size": 16,
  "raid_type": "raid_dp",
  "simulated_raid_groups": {
    "raid_type": "raid_dp"
  }
},
"storage_type": "hdd"
},
"cloud_storage": {
  "stores": {
    "cloud_store": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "store1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "used": 0
  }
},
"create_time": "2018-01-01 16:00:00 +0000",
"dr_home_node": {
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"home_node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"inactive_data_reporting": {
  "start_time": "2019-12-12 16:00:00 +0000"
},
"inode_attributes": {
  "file_private_capacity": 31136,

```

```

    "file_public_capacity": 31136,
    "files_private_used": 502,
    "files_total": 31136,
    "files_used": 97,
    "max_files_available": 31136,
    "max_files_possible": 2844525,
    "max_files_used": 97,
    "used_percent": 5,
    "version": 4
  },
  "is_spare_low": "",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"name": "node1",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},

```

```

"recommendation_spares": {
  "checksum_style": "block",
  "disk_class": "solid_state",
  "disk_type": "fc",
  "is_partition": 1,
  "layout_requirements": {
    "aggregate_min_disks": 6,
    "raid_group": {
      "default": 16,
      "max": 28,
      "min": 5
    },
    "raid_type": "raid_dp"
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "size": 10156769280,
  "syncmirror_pool": "pool0",
  "total": 10,
  "usable": 9
},
"snaplock_type": "non_snaplock",
"snapshot": {
  "files_total": 10,
  "files_used": 3,
  "max_files_available": 5,
  "max_files_used": 50
},
"space": {
  "block_storage": {
    "aggregate_metadata": 2655,
    "aggregate_metadata_percent": 8,
    "available": 10156560384,
    "data_compacted_count": 1990000,
    "data_compaction_space_saved": 1996000,
    "data_compaction_space_saved_percent": 27,
    "full_threshold_percent": 0,
    "inactive_user_data": 304448,
    "inactive_user_data_percent": 0,

```

```

    "physical_used": 2461696,
    "physical_used_percent": 50,
    "size": 10156769280,
    "used": 2088960,
    "used_including_snapshot_reserve": 674685,
    "used_including_snapshot_reserve_percent": 35,
    "used_percent": 50,
    "volume_deduplication_shared_count": 1990000,
    "volume_deduplication_space_saved": 1996000,
    "volume_deduplication_space_saved_percent": 27,
    "volume_footprints_percent": 14
  },
  "cloud_storage": {
    "used": 402743264
  },
  "efficiency": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots_flexclones": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "footprint": 608896,
  "snapshot": {
    "available": 2000,
    "reserve_percent": 20,
    "total": 5000,
    "used": 3000,
    "used_percent": 45
  }
},
"state": "online",
"statistics": {
  "iops_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  }
},

```

```

"latency_raw": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"status": "ok",
"throughput_raw": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 +0000"
},
"uuid": "string",
"volume-count": 0
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
787092	The target field cannot be specified for this operation.
7209049	Cannot perform the operation because the aggregate is currently expanding.
8586225	Unexpected error encountered when retrieving metrics and statistics for this aggregate.
19726382	Another provisioning operation is in progress on this cluster. Wait a few minutes, and try the operation again.
19726390	Unable to provide a recommendation to expand the aggregate.
19726391	Too many unassigned disks visible to the node that owns this aggregate.
19726392	Layout of this aggregate is not a supported configuration.
19726393	Failed to expand the aggregate. Aggregate expansion is not supported on this system.
19726394	Automatic aggregate expansion is not supported on systems with multiple data aggregates.

Error Code	Description
19726395	Automatic aggregate expansion is not supported when MetroCluster is not configured
19726396	Automatic aggregate expansion is not supported when the DR group is not in a normal state
19726397	Aggregates must contain disks with identical disk-types and disk-sizes.
19726402	Internal error. Unable to determine the MetroCluster configuration state.
19726538	Cannot perform the operation because the aggregate is not in a healthy state.

Name	Type	Description
error	error	

Example error

```

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

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

simulated_raid_groups

Name	Type	Description
added_data_disk_count	integer	Number of added data disks in RAID group.
added_parity_disk_count	integer	Number of added parity disks in RAID group.
existing_data_disk_count	integer	Number of existing data disks in the RAID group.
existing_parity_disk_count	integer	Number of existing parity disks in the RAID group.
is_partition	boolean	Indicates whether the disk is partitioned (true) or whole (false).
name	string	Name of the raid group.
usable_size	integer	Usable size of each disk, in bytes.

storage_pool_reference

Shared Storage Pool

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pools

Name	Type	Description
allocation_units_count	integer	Allocation count of storage pool.
storage_pool	storage_pool_reference	Shared Storage Pool

hybrid_cache

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

Name	Type	Description
disk_count	integer	Number of disks used in the cache tier of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
disk_type	string	Type of disk being used by the aggregate's cache tier.
enabled	boolean	Specifies whether the aggregate uses HDDs with SSDs as a cache.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
simulated_raid_groups	array[simulated_raid_groups]	
size	integer	Total usable space in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.
storage_pools	array[storage_pools]	List of storage pool properties and allocation_units_count for aggregate.
used	integer	Space used in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex_reference

Plex

Name	Type	Description
_links	_links	
name	string	

simulated_raid_groups

Name	Type	Description
added_data_disk_count	integer	Number of added data disks in RAID group.
added_parity_disk_count	integer	Number of added parity disks in RAID group.
data_disk_count	integer	Number of data disks in RAID group.
existing_data_disk_count	integer	Number of existing data disks in the RAID group.
existing_parity_disk_count	integer	Number of existing parity disks in the RAID group.
is_partition	boolean	Indicates whether the disk is partitioned (true) or whole (false).
name	string	Name of the raid group.
parity_disk_count	integer	Number of parity disks in RAID group.
raid_type	string	RAID type of the aggregate.
usable_size	integer	Usable size of each disk, in bytes.

primary

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

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

block_storage

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

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

Name	Type	Description
uses_partitions	boolean	If true, aggregate is using shared disks.

cloud_store

Cloud store

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cloud_storage_tier

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

cloud_storage

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

Name	Type	Description
attach_eligible	boolean	Specifies whether the aggregate is eligible for a cloud store to be attached.
migrate_threshold	integer	Specifies the minimum percentage of performance tier free space that must exist in order for migration of data from the capacity tier to performance tier to be allowed. Only valid for PATCH operations.
stores	array[cloud_storage_tier]	Configuration information for each cloud storage portion of the aggregate.

Name	Type	Description
tiering_fullness_threshold	integer	The percentage of space in the performance tier that must be used before data is tiered out to the cloud store. Only valid for PATCH operations.

data_encryption

Name	Type	Description
drive_protection_enabled	boolean	Specifies whether the aggregate uses self-encrypting drives with data protection enabled.
software_encryption_enabled	boolean	Specifies whether NetApp aggregate encryption is enabled. All data in the aggregate is encrypted.

dr_home_node

Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.

Name	Type	Description
name	string	
uuid	string	

home_node

Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

inactive_data_reporting

Name	Type	Description
enabled	boolean	Specifies whether or not inactive data reporting is enabled on the aggregate.

Name	Type	Description
start_time	string	Timestamp at which inactive data reporting was enabled on the aggregate.

inode_attributes

Name	Type	Description
file_private_capacity	integer	Number of files that can currently be stored on disk for system metadata files. This number will dynamically increase as more system files are created. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .
file_public_capacity	integer	Number of files that can currently be stored on disk for user-visible files. This number will dynamically increase as more user-visible files are created. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .
files_private_used	integer	Number of system metadata files used. If the referenced file system is restricted or offline, a value of 0 is returned. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .

Name	Type	Description
files_total	integer	Maximum number of user-visible files that this referenced file system can currently hold. If the referenced file system is restricted or offline, a value of 0 is returned.
files_used	integer	Number of user-visible files used in the referenced file system. If the referenced file system is restricted or offline, a value of 0 is returned.
max_files_available	integer	The count of the maximum number of user-visible files currently allowable on the referenced file system.
max_files_possible	integer	The largest value to which the maxfiles-available parameter can be increased by reconfiguration, on the referenced file system.
max_files_used	integer	The number of user-visible files currently in use on the referenced file system.
used_percent	integer	The percentage of disk space currently in use based on user-visible file count on the referenced file system.
version	integer	The inofile-version of the aggregate. If the referenced file system is restricted or offline, a value of 0 is returned. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .

iops

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

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

latency

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

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

throughput

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

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

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

metric

The most recent sample of I/O metrics for the aggregate.

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

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.

node

Node where the aggregate currently resides.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

raid_group

Name	Type	Description
default	integer	Default number of disks in a RAID group.

Name	Type	Description
max	integer	Maximum number of disks allowed in a RAID group.
min	integer	Minimum number of disks allowed in a RAID group.

layout_requirement

Name	Type	Description
aggregate_min_disks	integer	Minimum number of disks to create an aggregate.
default	boolean	Indicates if this RAID type is the default.
raid_group	raid_group	
raid_type	string	RAID type.

node

Node where the spares are assigned.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

aggregate_spare

Name	Type	Description
checksum_style	string	The checksum type that has been assigned to the spares.
disk_class	string	Disk class of spares.
disk_type	string	Type of disk.
is_partition	boolean	Indicates whether a disk is partitioned (true) or whole (false)
layout_requirements	array[layout_requirement]	Available RAID protections and their restrictions.

Name	Type	Description
node	node	Node where the spares are assigned.
size	integer	Usable size of each spare, in bytes.
syncmirror_pool	string	SyncMirror spare pool.
total	integer	Total number of spares in the bucket. The total spare count for each class of spares also includes reserved spare capacity recommended by ONTAP best practices. <ul style="list-style-type: none"> • example: 10 • readOnly: 1 • Introduced in: 9.11 • x-nullable: true
usable	integer	Total number of usable spares in the bucket. The usable count for each class of spares does not include reserved spare capacity recommended by ONTAP best practices. <ul style="list-style-type: none"> • example: 9 • readOnly: 1 • Introduced in: 9.6 • x-nullable: true

snapshot

Name	Type	Description
files_total	integer	Total files allowed in Snapshot copies
files_used	integer	Total files created in Snapshot copies
max_files_available	integer	Maximum files available for Snapshot copies

Name	Type	Description
max_files_used	integer	Files in use by Snapshot copies

block_storage

Name	Type	Description
aggregate_metadata	integer	Space used by different metafiles and internal operations inside the aggregate, in bytes.
aggregate_metadata_percent	integer	Aggregate metadata as a percentage.
available	integer	Space available in bytes.
data_compacted_count	integer	Amount of compacted data in bytes.
data_compaction_space_saved	integer	Space saved in bytes by compacting the data.
data_compaction_space_saved_percent	integer	Percentage saved by compacting the data.
full_threshold_percent	integer	The aggregate used percentage at which 'monitor.volume.full' EMS is generated.
inactive_user_data	integer	The size that is physically used in the block storage and has a cold temperature, in bytes. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>block_storage.inactive_user_data</code> or <code>**</code> .

Name	Type	Description
inactive_user_data_percent	integer	The percentage of inactive user data in the block storage. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>block_storage.inactive_user_data_percent</code> or <code>**</code> .
physical_used	integer	Total physical used size of an aggregate in bytes.
physical_used_percent	integer	Physical used percentage.
size	integer	Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
used	integer	Space used or reserved in bytes. Includes volume guarantees and aggregate metadata.
used_including_snapshot_reserve	integer	Total used including the Snapshot copy reserve, in bytes.
used_including_snapshot_reserve_percent	integer	Total used including the Snapshot reserve as a percentage.
used_percent	integer	Aggregate used percentage.
volume_deduplication_shared_count	integer	Amount of shared bytes counted by storage efficiency.
volume_deduplication_space_saved	integer	Amount of space saved in bytes by storage efficiency.
volume_deduplication_space_saved_percent	integer	Percentage of space saved by storage efficiency.

Name	Type	Description
volume_footprints_percent	integer	A summation of volume footprints inside the aggregate, as a percentage. A volume's footprint is the amount of space being used for the volume in the aggregate.

cloud_storage

Name	Type	Description
used	integer	Used space in bytes in the cloud store. Only applicable for aggregates with a cloud store tier.

efficiency

Storage efficiency.

Name	Type	Description
auto_adaptive_compression_savings	boolean	Indicates whether or not aggregate has auto adaptive compression savings.
cross_volume_background_dedupe	boolean	Indicates whether or not cross volume background deduplication is enabled.
cross_volume_dedupe_savings	boolean	Indicates whether or not aggregate has cross volume deduplication savings.
cross_volume_inline_dedupe	boolean	Indicates whether or not cross volume inline deduplication is enabled.
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

efficiency_without_snapshots

Storage efficiency that does not include the savings provided by Snapshot copies.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

efficiency_without_snapshots_flexclones

Storage efficiency that does not include the savings provided by Snapshot copies and Flexclone volumes.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

snapshot

Name	Type	Description
available	integer	Available space for Snapshot copies in bytes
reserve_percent	integer	Percentage of space reserved for Snapshot copies
total	integer	Total space for Snapshot copies in bytes
used	integer	Space used by Snapshot copies in bytes
used_percent	integer	Percentage of disk space used by Snapshot copies

space

Name	Type	Description
block_storage	block_storage	

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

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.

Name	Type	Description
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.

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

The real time I/O statistics for the aggregate.

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.
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.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Update an aggregate specified by the UUID

PATCH /storage/aggregates/{uuid}

Introduced In: 9.6

Updates the aggregate specified by the UUID with the properties in the body. This request starts a job and returns a link to that job.

Related ONTAP commands

- `storage aggregate add-disks`
- `storage aggregate mirror`
- `storage aggregate modify`
- `storage aggregate relocation start`
- `storage aggregate rename`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Aggregate UUID

Name	Type	In	Required	Description
auto_provision_policy	string	query	False	<p>If set to expand, the PATCH operation runs the recommended expansion of the aggregate.</p> <ul style="list-style-type: none"> • Introduced in: 9.8
simulate	boolean	query	False	<p>If set to true, the PATCH operation runs a simulated aggregate expansion with the provided input disk count and returns the proposed size of the new aggregate along with any associated warnings.</p> <ul style="list-style-type: none"> • Introduced in: 9.8
disk_size	integer	query	False	<p>If set, PATCH only selects disks of the specified size.</p>

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

Request Body

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

Name	Type	Description
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation.
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	home_node	Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.
inactive_data_reporting	inactive_data_reporting	
inode_attributes	inode_attributes	
is_spare_low	boolean	Specifies whether the aggregate is in a spares low condition on any of the RAID groups. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or **.
metric	metric	The most recent sample of I/O metrics for the aggregate.
name	string	Aggregate name.
node	node	Node where the aggregate currently resides.
recommendation_spares	array[aggregate_spare]	Information on the aggregate's remaining hot spare disks.
sidl_enabled	boolean	Specifies whether or not SIDL is enabled on the aggregate.

Name	Type	Description
snaplock_type	string	SnapLock type.
snapshot	snapshot	
space	space	
state	string	Operational state of the aggregate.
statistics	statistics	The real time I/O statistics for the aggregate.
uuid	string	Aggregate UUID.
volume-count	integer	Number of volumes in the aggregate.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "disk_type": "fc",
      "raid_size": 24,
      "raid_type": "raid_dp",
      "simulated_raid_groups": {
      },
      "size": 1612709888,
      "storage_pools": {
        "storage_pool": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "storage_pool_1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "used": 26501122
    },
    "mirror": {
      "enabled": "",
      "state": "unmirrored"
    },
    "plexes": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "plex0"
    }
  },
}
```

```

"primary": {
  "checksum_style": "block",
  "disk_class": "performance",
  "disk_count": 8,
  "disk_type": "fc",
  "raid_size": 16,
  "raid_type": "raid_dp",
  "simulated_raid_groups": {
    "raid_type": "raid_dp"
  }
},
"storage_type": "hdd"
},
"cloud_storage": {
  "stores": {
    "cloud_store": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "store1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "used": 0
  }
},
"create_time": "2018-01-01 16:00:00 +0000",
"dr_home_node": {
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"home_node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"inactive_data_reporting": {
  "start_time": "2019-12-12 16:00:00 +0000"
},
"inode_attributes": {
  "file_private_capacity": 31136,

```

```

    "file_public_capacity": 31136,
    "files_private_used": 502,
    "files_total": 31136,
    "files_used": 97,
    "max_files_available": 31136,
    "max_files_possible": 2844525,
    "max_files_used": 97,
    "used_percent": 5,
    "version": 4
  },
  "is_spare_low": "",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"name": "node1",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},

```

```

"recommendation_spares": {
  "checksum_style": "block",
  "disk_class": "solid_state",
  "disk_type": "fc",
  "is_partition": 1,
  "layout_requirements": {
    "aggregate_min_disks": 6,
    "raid_group": {
      "default": 16,
      "max": 28,
      "min": 5
    },
    "raid_type": "raid_dp"
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "size": 10156769280,
  "syncmirror_pool": "pool0",
  "total": 10,
  "usable": 9
},
"snaplock_type": "non_snaplock",
"snapshot": {
  "files_total": 10,
  "files_used": 3,
  "max_files_available": 5,
  "max_files_used": 50
},
"space": {
  "block_storage": {
    "aggregate_metadata": 2655,
    "aggregate_metadata_percent": 8,
    "available": 10156560384,
    "data_compacted_count": 1990000,
    "data_compaction_space_saved": 1996000,
    "data_compaction_space_saved_percent": 27,
    "full_threshold_percent": 0,
    "inactive_user_data": 304448,
    "inactive_user_data_percent": 0,

```

```

    "physical_used": 2461696,
    "physical_used_percent": 50,
    "size": 10156769280,
    "used": 2088960,
    "used_including_snapshot_reserve": 674685,
    "used_including_snapshot_reserve_percent": 35,
    "used_percent": 50,
    "volume_deduplication_shared_count": 1990000,
    "volume_deduplication_space_saved": 1996000,
    "volume_deduplication_space_saved_percent": 27,
    "volume_footprints_percent": 14
  },
  "cloud_storage": {
    "used": 402743264
  },
  "efficiency": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots_flexclones": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "footprint": 608896,
  "snapshot": {
    "available": 2000,
    "reserve_percent": 20,
    "total": 5000,
    "used": 3000,
    "used_percent": 45
  }
},
"state": "online",
"statistics": {
  "iops_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  }
},

```

```

"latency_raw": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"status": "ok",
"throughput_raw": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 +0000"
},
"uuid": "string",
"volume-count": 0
}

```

Response

Status: 200, Ok

Name	Type	Description
job	job_link	
num_records	integer	Number of records
records	array[aggregate]	
warnings	array[aggregate_warning]	List of validation warnings and remediation advice for the aggregate simulate behavior.

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "block_storage": {
    "hybrid_cache": {
      "disk_count": 6,
      "disk_type": "fc",
      "raid_size": 24,
      "raid_type": "raid_dp",
      "simulated_raid_groups": {
      },
      "size": 1612709888,
      "storage_pools": {
        "storage_pool": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "storage_pool_1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "used": 26501122
    },
    "mirror": {
      "enabled": "",

```



```

    "state": "unmirrored"
  },
  "plexes": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "plex0"
  },
  "primary": {
    "checksum_style": "block",
    "disk_class": "performance",
    "disk_count": 8,
    "disk_type": "fc",
    "raid_size": 16,
    "raid_type": "raid_dp",
    "simulated_raid_groups": {
      "raid_type": "raid_dp"
    }
  },
  "storage_type": "hdd"
},
"cloud_storage": {
  "stores": {
    "cloud_store": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "store1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "used": 0
  }
},
"create_time": "2018-01-01 16:00:00 +0000",
"dr_home_node": {
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"home_node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
}

```

```

    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"inactive_data_reporting": {
  "start_time": "2019-12-12 16:00:00 +0000"
},
"inode_attributes": {
  "file_private_capacity": 31136,
  "file_public_capacity": 31136,
  "files_private_used": 502,
  "files_total": 31136,
  "files_used": 97,
  "max_files_available": 31136,
  "max_files_possible": 2844525,
  "max_files_used": 97,
  "used_percent": 5,
  "version": 4
},
"is_spare_low": "",
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"duration": "PT15S",
"iops": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"latency": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"status": "ok",
"throughput": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 +0000"
},

```

```

"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"recommendation_spare": {
  "checksum_style": "block",
  "disk_class": "solid_state",
  "disk_type": "fc",
  "is_partition": 1,
  "layout_requirements": {
    "aggregate_min_disks": 6,
    "raid_group": {
      "default": 16,
      "max": 28,
      "min": 5
    }
  },
  "raid_type": "raid_dp"
},
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": 10156769280,
"syncmirror_pool": "pool0",
"total": 10,
"usable": 9
},
"snaplock_type": "non_snaplock",
"snapshot": {
  "files_total": 10,
  "files_used": 3,
  "max_files_available": 5,
  "max_files_used": 50
},
"space": {

```

```

"block_storage": {
  "aggregate_metadata": 2655,
  "aggregate_metadata_percent": 8,
  "available": 10156560384,
  "data_compacted_count": 1990000,
  "data_compaction_space_saved": 1996000,
  "data_compaction_space_saved_percent": 27,
  "full_threshold_percent": 0,
  "inactive_user_data": 304448,
  "inactive_user_data_percent": 0,
  "physical_used": 2461696,
  "physical_used_percent": 50,
  "size": 10156769280,
  "used": 2088960,
  "used_including_snapshot_reserve": 674685,
  "used_including_snapshot_reserve_percent": 35,
  "used_percent": 50,
  "volume_deduplication_shared_count": 1990000,
  "volume_deduplication_space_saved": 1996000,
  "volume_deduplication_space_saved_percent": 27,
  "volume_footprints_percent": 14
},
"cloud_storage": {
  "used": 402743264
},
"efficiency": {
  "logical_used": 0,
  "ratio": 0,
  "savings": 0
},
"efficiency_without_snapshots": {
  "logical_used": 0,
  "ratio": 0,
  "savings": 0
},
"efficiency_without_snapshots_flexclones": {
  "logical_used": 0,
  "ratio": 0,
  "savings": 0
},
"footprint": 608896,
"snapshot": {
  "available": 2000,
  "reserve_percent": 20,
  "total": 5000,
  "used": 3000,

```

```
    "used_percent": 45
  },
  "state": "online",
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "uuid": "string",
  "volume-count": 0
},
"warnings": {
  "action": {
    "arguments": {
    }
  },
  "warning": {
    "arguments": {
    }
  }
}
}
```

Response

Status: 202, Accepted

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
262247	The value is invalid for the field.
460777	Failed to get information on the job.
786434	Cannot connect to node where the aggregate resides.
786435	Internal Error. Failed to create a communication handle.
786439	An aggregate already uses the specified name.
786447	Failed to modify the aggregate.
786456	Failed to add disks to the aggregate.
786458	Failed to rename aggregate.
786465	Failed to take aggregate offline because of the reason outlined in the message.
786467	Failed to take aggregate online because of the reason outlined in the message.
786468	VLDB is offline.
786472	Node that hosts the aggregate is offline.
786479	Cannot find node ID for the node.
786491	Not enough spares on the node.
786730	Internal Error
786771	Aggregate does not exist.
786787	Aggregate is not online.
786808	Aggregate mirror failed.
786867	Specified aggregate resides on the remote cluster.
786911	Not every node in the cluster has the Data ONTAP version required for the feature.
786923	This operation is not allowed during the pre-commit phase of a 7-mode to clustered Data ONTAP transition.
786924	Internal error for an aggregate that is in the pre-commit phase of a 7-mode to clustered Data ONTAP transition.
786955	Modifying raidtype to raid_tec requires a minimum of six disks in the RAID Group.

Error Code	Description
786956	Modifying raidtype to raid_dp requires a minimum of four disks in the RAID Group.
786965	Spare Selection in userspace failed.
787046	Mirroring of a FabricPool is not allowed.
787092	The target field cannot be specified for this operation.
787144	Aggregate is not a FabricPool.
787156	Modifying the attributes of mirror object store is not allowed.
787169	Only one field can be modified per operation.
787170	Failed to patch the "block_storage.primary.disk_count" because the disk count specified is smaller than existing disk count.
787172	This query is only allowed during the modification of the specified field.
787178	Unmirroring an aggregate with a PATCH operation is not supported.
787187	Internal error. Failed to check if the aggregate is a FabricPool.
787266	Invalid aggregate state. This state is not supported for a PATCH operation.
787273	Allocation unit count is not valid.
787274	Raidtype is not valid.
787275	Patch request with multiple records is not valid.
787276	Storage pool name and uuid do not match.
787277	Storage pool name and uuid are empty.
787278	Incorrect storage pool name specified.
787279	Incorrect storage pool uuid specified.
787280	Allocation unit count specified is smaller than the existing allocation unit count.
787281	Cannot modify RAID type of aggregate hybrid cache tier.
787282	RAID group must be specified on a disk addition to an aggregate with mixed RAID types.
787283	RAID group must be specified on a disk addition to a Flash Pool aggregate.
787284	The specified RAID group uses capacity from one or more storage pools.

Error Code	Description
787287	Cannot add physical SSD cache because aggregate uses cache capacity from a storage pool.
787288	Cannot add storage pool units because aggregate uses physical SSD cache.
787289	Incorrect raid_group specified during first time addition of physical SSD cache to an aggregate.
787291	Cannot specify RAID group which is located on the aggregate primary tier.
787293	Cannot specify RAID group which is located on the aggregate cache tier.
787294	This query is only allowed during the modification of the specific fields.
787295	The storage pool allocation units count is required.
1258699	Cannot use all the disks specified for the requested operation.
1263500	Operation will lead to creation of new raid group.
1263501	Operation will exceed half of the maximum volume sizes allowed on the node.
1263502	One spare data partition from at least one of the chosen root-data1-data2 disks will not be used.
1263503	Operation will lead to downsizing of one or more disks.
1263504	Operation will lead to a spares low condition.
1263598	One or more selected disks will be partitioned.
1263624	Operation will lead to a no sparecore condition.
2425736	No matching node found for the UUID provided.
7208962	Aggregate in an inconsistent state.
7208993	Failed to offline as the volume is being used.
7209033	CIFS open files prevent operation.
7209049	Cannot perform the operation because the aggregate is currently expanding.
7209075	Cannot perform the operation because the volume size limit for this system type would be exceeded.
7209090	Inconsistent state.
7209183	Volume is a partial volume.
7209229	This version of ONTAP does not recognize the filesystem. It is probably from a later version of the software and is being left offline.

Error Code	Description
7209246	The specified operation could not be completed as the volume is currently busy.
7209247	The volume was not found.
7209263	Container has failed.
7209271	wafiron is currently active.
7209275	Container was created in an unclustered ONTAP deployment.
7209463	Nvfile replay pending.
7209966	Another online request is already in progress for aggregate. The previous online request is waiting on a response from the licensing manager.
11206666	Storage pool is not healthy.
11210659	Aggregate is not online.
11210662	Adding capacity from storage pool to a mirrored aggregate is not supported.
11210667	Storage pool does not have enough spare allocation units.
11210670	Cannot add capacity from storage pool to aggregate, because currently allocated capacity to the aggregate does not span across all disks belonging to the storage pool.
11210672	Cannot grow aggregate as no capacity is allocated to it from storage pool.
11210673	Mixing of physical SSDs and capacity from a storage pool is not allowed in same aggregate.
11210675	Capacity in storage pool belongs to different fault isolation domain than aggregate.
11210678	Storage pool does not have enough disks to create RAID groups of same raid type as that of already allocated cache tier.
11210679	Storage pool does not have enough disks to create RAID groups of type RAID-DP.
11210680	Storage pool does not have enough disks to create RAID groups of same raid type as that of the aggregate.
11210685	Storage pool does not have enough disks to create RAID groups of type RAID-TEC.
11210688	Capacity from storage pool cannot be added to an SSD aggregate and a Flash Pool.

Error Code	Description
13108106	Cannot run aggregate relocation because volume expand is in progress.
19726347	There are a number of unassigned disks visible to the node that owns this aggregate.
19726382	Another provisioning operation is in progress on this cluster. Wait a few minutes, and try the operation again.
19726390	Unable to automatically expand this aggregate.
19726391	Too many unassigned disks visible to the node that owns this aggregate.
19726392	Layout of this aggregate is not a supported configuration.
19726393	Failed to expand the aggregate. Aggregate expansion is not supported on this system.
19726394	Automatic aggregate expansion is not supported on systems with multiple data aggregates.
19726395	Automatic aggregate expansion is not supported when MetroCluster is not configured.
19726396	Automatic aggregate expansion is not supported when the DR group is not in a normal state.
19726397	Aggregates must contain disks with identical disk-types and disk-sizes.
19726402	Internal error. Unable to determine the MetroCluster configuration state.
19726538	Cannot perform the operation because the aggregate is not in a healthy state.
26542083	Destination node is at higher Data ONTAP version than source node.
26542084	Source node is at higher Data ONTAP version than destination node.
26542097	Unable to get D-blade ID of destination.
26542101	Unable to contact the source node.
26542102	Unable to contact the destination node.
26542120	An SVM migrate operation is in progress. When the migrate operation completes, try the operation again.
26542121	A MetroCluster disaster recovery operation is in progress. When the recovery operation completes, try the operation again.
196608334	Failed to modify the aggregate because it contains NAE volumes.

Error Code	Description
196608335	Failed to modify the aggregate because it contains non-encrypted volumes.

Name	Type	Description
error	error	

Example error

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

simulated_raid_groups

Name	Type	Description
added_data_disk_count	integer	Number of added data disks in RAID group.
added_parity_disk_count	integer	Number of added parity disks in RAID group.
existing_data_disk_count	integer	Number of existing data disks in the RAID group.
existing_parity_disk_count	integer	Number of existing parity disks in the RAID group.
is_partition	boolean	Indicates whether the disk is partitioned (true) or whole (false).
name	string	Name of the raid group.
usable_size	integer	Usable size of each disk, in bytes.

storage_pool_reference

Shared Storage Pool

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pools

Name	Type	Description
allocation_units_count	integer	Allocation count of storage pool.
storage_pool	storage_pool_reference	Shared Storage Pool

hybrid_cache

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

Name	Type	Description
disk_count	integer	Number of disks used in the cache tier of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
disk_type	string	Type of disk being used by the aggregate's cache tier.
enabled	boolean	Specifies whether the aggregate uses HDDs with SSDs as a cache.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
simulated_raid_groups	array[simulated_raid_groups]	
size	integer	Total usable space in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.
storage_pools	array[storage_pools]	List of storage pool properties and allocation_units_count for aggregate.
used	integer	Space used in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex_reference

Plex

Name	Type	Description
_links	_links	
name	string	

simulated_raid_groups

Name	Type	Description
added_data_disk_count	integer	Number of added data disks in RAID group.
added_parity_disk_count	integer	Number of added parity disks in RAID group.
data_disk_count	integer	Number of data disks in RAID group.
existing_data_disk_count	integer	Number of existing data disks in the RAID group.
existing_parity_disk_count	integer	Number of existing parity disks in the RAID group.
is_partition	boolean	Indicates whether the disk is partitioned (true) or whole (false).
name	string	Name of the raid group.
parity_disk_count	integer	Number of parity disks in RAID group.
raid_type	string	RAID type of the aggregate.
usable_size	integer	Usable size of each disk, in bytes.

primary

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

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

block_storage

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

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

Name	Type	Description
uses_partitions	boolean	If true, aggregate is using shared disks.

cloud_store

Cloud store

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cloud_storage_tier

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

cloud_storage

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

Name	Type	Description
attach_eligible	boolean	Specifies whether the aggregate is eligible for a cloud store to be attached.
migrate_threshold	integer	Specifies the minimum percentage of performance tier free space that must exist in order for migration of data from the capacity tier to performance tier to be allowed. Only valid for PATCH operations.
stores	array[cloud_storage_tier]	Configuration information for each cloud storage portion of the aggregate.

Name	Type	Description
tiering_fullness_threshold	integer	The percentage of space in the performance tier that must be used before data is tiered out to the cloud store. Only valid for PATCH operations.

data_encryption

Name	Type	Description
drive_protection_enabled	boolean	Specifies whether the aggregate uses self-encrypting drives with data protection enabled.
software_encryption_enabled	boolean	Specifies whether NetApp aggregate encryption is enabled. All data in the aggregate is encrypted.

dr_home_node

Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.

Name	Type	Description
name	string	
uuid	string	

home_node

Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

inactive_data_reporting

Name	Type	Description
enabled	boolean	Specifies whether or not inactive data reporting is enabled on the aggregate.

Name	Type	Description
start_time	string	Timestamp at which inactive data reporting was enabled on the aggregate.

inode_attributes

Name	Type	Description
file_private_capacity	integer	Number of files that can currently be stored on disk for system metadata files. This number will dynamically increase as more system files are created. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .
file_public_capacity	integer	Number of files that can currently be stored on disk for user-visible files. This number will dynamically increase as more user-visible files are created. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .
files_private_used	integer	Number of system metadata files used. If the referenced file system is restricted or offline, a value of 0 is returned. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .

Name	Type	Description
files_total	integer	Maximum number of user-visible files that this referenced file system can currently hold. If the referenced file system is restricted or offline, a value of 0 is returned.
files_used	integer	Number of user-visible files used in the referenced file system. If the referenced file system is restricted or offline, a value of 0 is returned.
max_files_available	integer	The count of the maximum number of user-visible files currently allowable on the referenced file system.
max_files_possible	integer	The largest value to which the maxfiles-available parameter can be increased by reconfiguration, on the referenced file system.
max_files_used	integer	The number of user-visible files currently in use on the referenced file system.
used_percent	integer	The percentage of disk space currently in use based on user-visible file count on the referenced file system.
version	integer	The inofile-version of the aggregate. If the referenced file system is restricted or offline, a value of 0 is returned. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** .

iops

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

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

latency

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

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

throughput

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

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

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

metric

The most recent sample of I/O metrics for the aggregate.

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

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.

node

Node where the aggregate currently resides.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

raid_group

Name	Type	Description
default	integer	Default number of disks in a RAID group.

Name	Type	Description
max	integer	Maximum number of disks allowed in a RAID group.
min	integer	Minimum number of disks allowed in a RAID group.

layout_requirement

Name	Type	Description
aggregate_min_disks	integer	Minimum number of disks to create an aggregate.
default	boolean	Indicates if this RAID type is the default.
raid_group	raid_group	
raid_type	string	RAID type.

node

Node where the spares are assigned.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

aggregate_spare

Name	Type	Description
checksum_style	string	The checksum type that has been assigned to the spares.
disk_class	string	Disk class of spares.
disk_type	string	Type of disk.
is_partition	boolean	Indicates whether a disk is partitioned (true) or whole (false)
layout_requirements	array[layout_requirement]	Available RAID protections and their restrictions.

Name	Type	Description
node	node	Node where the spares are assigned.
size	integer	Usable size of each spare, in bytes.
syncmirror_pool	string	SyncMirror spare pool.
total	integer	Total number of spares in the bucket. The total spare count for each class of spares also includes reserved spare capacity recommended by ONTAP best practices. <ul style="list-style-type: none"> • example: 10 • readOnly: 1 • Introduced in: 9.11 • x-nullable: true
usable	integer	Total number of usable spares in the bucket. The usable count for each class of spares does not include reserved spare capacity recommended by ONTAP best practices. <ul style="list-style-type: none"> • example: 9 • readOnly: 1 • Introduced in: 9.6 • x-nullable: true

snapshot

Name	Type	Description
files_total	integer	Total files allowed in Snapshot copies
files_used	integer	Total files created in Snapshot copies
max_files_available	integer	Maximum files available for Snapshot copies

Name	Type	Description
max_files_used	integer	Files in use by Snapshot copies

block_storage

Name	Type	Description
aggregate_metadata	integer	Space used by different metafiles and internal operations inside the aggregate, in bytes.
aggregate_metadata_percent	integer	Aggregate metadata as a percentage.
available	integer	Space available in bytes.
data_compacted_count	integer	Amount of compacted data in bytes.
data_compaction_space_saved	integer	Space saved in bytes by compacting the data.
data_compaction_space_saved_percent	integer	Percentage saved by compacting the data.
full_threshold_percent	integer	The aggregate used percentage at which 'monitor.volume.full' EMS is generated.
inactive_user_data	integer	The size that is physically used in the block storage and has a cold temperature, in bytes. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>block_storage.inactive_user_data</code> or <code>**</code> .

Name	Type	Description
inactive_user_data_percent	integer	The percentage of inactive user data in the block storage. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>block_storage.inactive_user_data_percent</code> or <code>**</code> .
physical_used	integer	Total physical used size of an aggregate in bytes.
physical_used_percent	integer	Physical used percentage.
size	integer	Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
used	integer	Space used or reserved in bytes. Includes volume guarantees and aggregate metadata.
used_including_snapshot_reserve	integer	Total used including the Snapshot copy reserve, in bytes.
used_including_snapshot_reserve_percent	integer	Total used including the Snapshot reserve as a percentage.
used_percent	integer	Aggregate used percentage.
volume_deduplication_shared_count	integer	Amount of shared bytes counted by storage efficiency.
volume_deduplication_space_saved	integer	Amount of space saved in bytes by storage efficiency.
volume_deduplication_space_saved_percent	integer	Percentage of space saved by storage efficiency.

Name	Type	Description
volume_footprints_percent	integer	A summation of volume footprints inside the aggregate, as a percentage. A volume's footprint is the amount of space being used for the volume in the aggregate.

cloud_storage

Name	Type	Description
used	integer	Used space in bytes in the cloud store. Only applicable for aggregates with a cloud store tier.

efficiency

Storage efficiency.

Name	Type	Description
auto_adaptive_compression_savings	boolean	Indicates whether or not aggregate has auto adaptive compression savings.
cross_volume_background_dedupe	boolean	Indicates whether or not cross volume background deduplication is enabled.
cross_volume_dedupe_savings	boolean	Indicates whether or not aggregate has cross volume deduplication savings.
cross_volume_inline_dedupe	boolean	Indicates whether or not cross volume inline deduplication is enabled.
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

efficiency_without_snapshots

Storage efficiency that does not include the savings provided by Snapshot copies.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

efficiency_without_snapshots_flexclones

Storage efficiency that does not include the savings provided by Snapshot copies and Flexclone volumes.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

snapshot

Name	Type	Description
available	integer	Available space for Snapshot copies in bytes
reserve_percent	integer	Percentage of space reserved for Snapshot copies
total	integer	Total space for Snapshot copies in bytes
used	integer	Space used by Snapshot copies in bytes
used_percent	integer	Percentage of disk space used by Snapshot copies

space

Name	Type	Description
block_storage	block_storage	

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

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.

Name	Type	Description
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.

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

The real time I/O statistics for the aggregate.

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.
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.

aggregate

Name	Type	Description
<code>_links</code>	_links	
<code>_tags</code>	array[string]	Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings.
<code>block_storage</code>	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
<code>cloud_storage</code>	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
<code>create_time</code>	string	Timestamp of aggregate creation.
<code>data_encryption</code>	data_encryption	
<code>dr_home_node</code>	dr_home_node	Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.
<code>home_node</code>	home_node	Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.
<code>inactive_data_reporting</code>	inactive_data_reporting	
<code>inode_attributes</code>	inode_attributes	
<code>is_spare_low</code>	boolean	Specifies whether the aggregate is in a spares low condition on any of the RAID groups. This is an advanced property; there is an added computational cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>footprint</code> or <code>**</code> .

Name	Type	Description
metric	metric	The most recent sample of I/O metrics for the aggregate.
name	string	Aggregate name.
node	node	Node where the aggregate currently resides.
recommendation_spares	array[aggregate_spare]	Information on the aggregate's remaining hot spare disks.
sidl_enabled	boolean	Specifies whether or not SIDL is enabled on the aggregate.
snaplock_type	string	SnapLock type.
snapshot	snapshot	
space	space	
state	string	Operational state of the aggregate.
statistics	statistics	The real time I/O statistics for the aggregate.
uuid	string	Aggregate UUID.
volume-count	integer	Number of volumes in the aggregate.

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.

action

Name	Type	Description
arguments	array[string]	Arguments present in the specified action message.

Name	Type	Description
code	integer	Corrective action code of the specified action.
message	string	Specifies the corrective action to be taken to resolve the issue.

warning

Name	Type	Description
arguments	array[string]	Arguments present in the warning message encountered.
code	integer	Warning code of the warning encountered.
message	string	Details of the warning encountered by the aggregate simulate query.

aggregate_warning

Name	Type	Description
action	action	
name	string	Name of the entity that returns the warning.
warning	warning	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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

Storage aggregate metrics

Storage aggregates UUID metrics endpoint overview

Overview

The Storage Aggregate Metrics API provides historical performance metrics for the specified aggregate. The collection GET operation retrieves read, write, other and total metrics for a given aggregate, in terms of IOPS, latency and throughput. The read and write categories display the I/O operations that service user reads and writes across all the hosted volumes on a given aggregate. The other category encompasses background I/O operations that implement data protection services currently running on the aggregate. IOPs are the number of I/O operations reported per second, throughput is the amount of I/O operations measured in bytes per second and latency is the average response time for an IOP, reported in microseconds. Without a specified time interval, the output is limited to statistics collected at 15 second intervals over the last hour.

Examples

Retrieving metrics for an aggregate

In this example, the API returns a set of records that exist for the aggregate with the given UUID for the last hour.

```
# The API:
/api/storage/aggregates/{uuid}/metrics

#The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/538bf337-1b2c-11e8-
bad0-005056b48388/metrics?max_records=4" -H "accept: application/json"

#The response:
{
  "records": [
    {
      "timestamp": "2019-01-14T23:33:45Z"
    },
    {
      "timestamp": "2019-01-14T23:33:30Z"
    },
    {
      "timestamp": "2019-01-14T23:33:15Z"
    },
    {
      "timestamp": "2019-01-14T23:33:00Z"
    }
  ],
  "num_records": 4
}
```

Retrieving metrics for an aggregate with a set timestamp

In this example, the API returns metric values for latency, IOPS, and throughput properties such as read, write and total. The status and duration for which the metrics are requested are also returned.

```

#The API:
/api/storage/aggregates/{uuid}/metrics?timestamp={timestamp}

#The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/538bf337-1b2c-11e8-
bad0-005056b48388/metrics?timestamp=2019-01-1T23:33:00Z" -H "accept:
application/json"

#The response:
{
  "records": [
    {
      "uuid": "538bf337-1b2c-11e8-bad0-005056b48388",
      "timestamp": "2019-01-01T23:33:00Z",
      "status": "ok",
      "duration": "PT15S",
      "throughput": {
        "read": 6826,
        "write": 205892,
        "other": 0,
        "total": 212718
      },
      "latency": {
        "read": 148,
        "write": 216,
        "other": 0,
        "total": 199
      },
      "iops": {
        "read": 1,
        "write": 5,
        "other": 0,
        "total": 6
      }
    }
  ]
}

```

Retrieving metrics for an aggregate for a set time interval

In this example, the API returns the requested metrics for the given time interval of 1 week. The interval value can be 1 hour, 1 day, 1 week, 1 month or 1 year. If the interval value is not set, a default value of 1 hour is used.

```
#The API:
```

```
/api/storage/aggregates/{uuid}/metrics
```

```
#The call:
```

```
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/538bf337-1b2c-11e8-bad0-005056b48388/metrics?return_timeout=15&fields=*&interval=1w&max_records=4" -H "accept: application/json"
```

```
#The response:
```

```
{
  "records": [
    {
      "timestamp": "2019-01-01T23:30:00Z",
      "status": "ok",
      "duration": "PT30M",
      "throughput": {
        "read": 268328,
        "write": 5556255,
        "other": 0,
        "total": 5824584
      },
      "latency": {
        "read": 156,
        "write": 430,
        "other": 0,
        "total": 318
      },
      "iops": {
        "read": 18,
        "write": 26,
        "other": 0,
        "total": 45
      }
    },
    {
      "timestamp": "2019-01-01T23:00:00Z",
      "status": "ok",
      "duration": "PT30M",
      "throughput": {
        "read": 474266,
        "write": 6121908,
        "other": 0,
        "total": 6596175
      },
      "latency": {
        "read": 154,
```

```
    "write": 448,  
    "other": 0,  
    "total": 262  
  },  
  "iops": {  
    "read": 48,  
    "write": 28,  
    "other": 0,  
    "total": 76  
  }  
},  
{  
  "timestamp": "2019-01-01T22:30:00Z",  
  "status": "ok",  
  "duration": "PT30M",  
  "throughput": {  
    "read": 540164,  
    "write": 2411356,  
    "other": 26244685,  
    "total": 29196206  
  },  
  "latency": {  
    "read": 159,  
    "write": 394,  
    "other": 192,  
    "total": 193  
  },  
  "iops": {  
    "read": 94,  
    "write": 16,  
    "other": 437,  
    "total": 548  
  }  
},  
{  
  "timestamp": "2019-01-01T22:00:00Z",  
  "status": "ok",  
  "duration": "PT30M",  
  "throughput": {  
    "read": 2842,  
    "write": 2765407,  
    "other": 0,  
    "total": 2768249  
  },  
  "latency": {  
    "read": 189,
```

```

    "write": 540,
    "other": 0,
    "total": 523
  },
  "iops": {
    "read": 0,
    "write": 13,
    "other": 0,
    "total": 13
  }
},
"num_records": 4
}

```

Related ONTAP commands

- `statistics aggregate show`

Retrieve historical performance metrics for an aggregate

GET `/storage/aggregates/{uuid}/metrics`

Introduced In: 9.7

Retrieves historical performance metrics for an aggregate.

Parameters

Name	Type	In	Required	Description
timestamp	string	query	False	Filter by timestamp
status	string	query	False	Filter by status
throughput.read	integer	query	False	Filter by throughput.read
throughput.other	integer	query	False	Filter by throughput.other
throughput.write	integer	query	False	Filter by throughput.write
throughput.total	integer	query	False	Filter by throughput.total

Name	Type	In	Required	Description
duration	string	query	False	Filter by duration
latency.read	integer	query	False	Filter by latency.read
latency.other	integer	query	False	Filter by latency.other
latency.write	integer	query	False	Filter by latency.write
latency.total	integer	query	False	Filter by latency.total
iops.read	integer	query	False	Filter by iops.read
iops.other	integer	query	False	Filter by iops.other
iops.write	integer	query	False	Filter by iops.write
iops.total	integer	query	False	Filter by iops.total
uuid	string	path	True	Unique identifier of the aggregate.

Name	Type	In	Required	Description
interval	string	query	False	<p>The time range for the data. Values can be 1h, 1d, 1w, 1m, or 1y. The period for each time range is as follows:</p> <ul style="list-style-type: none"> • 1h: Metrics over the most recent hour sampled over 15 seconds. • 1d: Metrics over the most recent day sampled over 5 minutes. • 1w: Metrics over the most recent week sampled over 30 minutes. • 1m: Metrics over the most recent month sampled over 2 hours. • 1y: Metrics over the most recent year sampled over a day. • Default value: 1 • enum: ["1h", "1d", "1w", "1m", "1y"]

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

Response

Status: 200, Ok

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
8586225	Encountered unexpected error in retrieving metrics for the requested aggregate.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

iops

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

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

latency

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

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

throughput

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

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

records

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage bridges

Storage bridges endpoint overview

Retrieving storage bridge information

The storage bridge GET API retrieves all of the bridges in the cluster.

Examples

1) Retrieves a list of bridges from the cluster

The following example shows the response with a list of bridges from the cluster:

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

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

# The response:
{
  "records": [
    {
      "wwn": "2000001086a18100",
      "name": "ATTO_2000001086a18100",
      "_links": {
        "self": {
          "href": "/api/storage/bridges/2000001086a18100"
        }
      }
    },
    {
      "wwn": "2000001086a18380",
      "name": "ATTO_2000001086a18380",
      "_links": {
        "self": {
          "href": "/api/storage/bridges/2000001086a18380"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/storage/bridges/"
    }
  }
}
```

2) Retrieves a specific bridge from the cluster

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



```

# The API:
/api/storage/bridges/{wwn}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/bridges/2000001086a18100" -H
"accept: application/hal+json"

# The response:
{
  "wwn": "2000001086a18100",
  "chassis_throughput_state": "ok",
  "dram_single_bit_error_count": 0,
  "firmware_version": "3.10 007A",
  "ip_address": "10.226.57.178",
  "security_enabled": false,
  "monitoring_enabled": true,
  "model": "FibreBridge 7500N",
  "state": "ok",
  "managed_by": "in_band",
  "serial_number": "FB7500N102450",
  "symbolic_name": "RTP-FCSAS02-41KK10",
  "vendor": "atto",
  "name": "ATTO_2000001086a18100",
  "last_reboot": {
    "reason": {
      "message": "Reason: \"FirmwareRestart Command\".",
      "code": "39321683"
    }
  },
  "time": "2020-12-09T00:47:58-05:00"
},
"paths": [
{
  "name": "0e",
  "node": {
    "name": "sti8080mcc-htp-005",
    "uuid": "ecc3d992-3a86-11eb-9fab-00a0985a6024",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/ecc3d992-3a86-11eb-9fab-00a0985a6024"
      }
    }
  }
},
  "_links": {
    "self": {
      "href": "/api/storage/ports/ecc3d992-3a86-11eb-9fab-
00a0985a6024/0e"
    }
  }
}

```

```

    }
  },
  "target_port": {
    "wwn": "2100001086a18380"
  }
}
],
"temperature_sensor": {
  "name": "Chassis Temperature Sensor",
  "minimum": 0,
  "maximum": 90,
  "reading": 54,
  "state": "ok"
},
"fc_ports": [
{
  "id": 1,
  "state": "online",
  "enabled": true,
  "negotiated_data_rate": 8,
  "configured_data_rate": 8,
  "conn_mode": "ptp",
  "data_rate_capability": 16,
  "wwn": "2100001086a18100",
  "peer_wwn": "0000000000000000",
  "sfp": {
    "vendor": "FINISAR CORP.",
    "serial_number": "UW106SA",
    "part_number": "FTLF8529P3BCV",
    "data_rate_capability": 16
  }
},
{
  "id": 2,
  "state": "online",
  "enabled": true,
  "negotiated_data_rate": 16,
  "configured_data_rate": 16,
  "conn_mode": "ptp",
  "data_rate_capability": 16,
  "wwn": "2200001086a18100",
  "peer_wwn": "0000000000000000",
  "sfp": {
    "vendor": "FINISAR CORP.",
    "serial_number": "UW1072B",
    "part_number": "FTLF8529P3BCV",

```

```

    "data_rate_capability": 16
  }
}
],
"power_supply_units": [
{
  "name": "A",
  "state": "ok"
},
{
  "name": "B",
  "state": "ok"
}
],
"sas_ports": [
{
  "id": 1,
  "state": "online",
  "enabled": true,
  "data_rate_capability": 12,
  "negotiated_data_rate": 6,
  "wwn": "5001086000a18100",
  "phy_1": {
    "state": "online"
  },
  "phy_2": {
    "state": "online"
  },
  "phy_3": {
    "state": "online"
  },
  "phy_4": {
    "state": "online"
  },
  "cable": {
    "vendor": "Molex Inc.",
    "serial_number": "618130935",
    "technology": "Passive Copper 5m ID:00",
    "part_number": "112-00431"
  }
},
{
  "state": "offline",
  "enabled": false,
  "data_rate_capability": 12,
  "negotiated_data_rate": 0,

```

```

"wwn": "5001086000a18104",
"phy_1": {
  "state": "offline"
},
"phy_2": {
  "state": "offline"
},
"phy_3": {
  "state": "offline"
},
"phy_4": {
  "state": "offline"
},
},
{
  "state": "offline",
  "enabled": false,
  "data_rate_capability": 12,
  "negotiated_data_rate": 0,
  "wwn": "5001086000a18108",
  "phy_1": {
    "state": "offline"
  },
  "phy_2": {
    "state": "offline"
  },
  "phy_3": {
    "state": "offline"
  },
  "phy_4": {
    "state": "offline"
  },
},
{
  "state": "offline",
  "enabled": false,
  "data_rate_capability": 12,
  "negotiated_data_rate": 0,
  "wwn": "5001086000a1810c",
  "phy_1": {
    "state": "offline"
  },
  "phy_2": {
    "state": "offline"
  },
  "phy_3": {

```

```

    "state": "offline"
  },
  "phy_4": {
    "state": "offline"
  },
}
],
"_links": {
"self": {
  "href": "/api/storage/bridges/2000001086a18100"
}
}
}

```

Retrieve a collection of bridges

GET /storage/bridges

Introduced In: 9.9

Retrieves a collection of bridges.

Related ONTAP commands

- `storage bridge show`

Learn more

- [DOC /storage/bridges](#)

Parameters

Name	Type	In	Required	Description
wwn	string	query	False	Filter by wwn
managed_by	string	query	False	Filter by managed_by
paths.node.uuid	string	query	False	Filter by paths.node.uuid
paths.node.name	string	query	False	Filter by paths.node.name

Name	Type	In	Required	Description
paths.source_port.id	string	query	False	Filter by paths.source_port.id
paths.source_port.name	string	query	False	Filter by paths.source_port.name
paths.target_port.id	string	query	False	Filter by paths.target_port.id
paths.target_port.name	string	query	False	Filter by paths.target_port.name
paths.target_port.wwn	string	query	False	Filter by paths.target_port.wwn
paths.name	string	query	False	Filter by paths.name
state	string	query	False	Filter by state
last_reboot.time	string	query	False	Filter by last_reboot.time
last_reboot.reason.arguments.message	string	query	False	Filter by last_reboot.reason.arguments.message • Introduced in: 9.10
last_reboot.reason.arguments.code	string	query	False	Filter by last_reboot.reason.arguments.code • Introduced in: 9.10
last_reboot.reason.message	string	query	False	Filter by last_reboot.reason.message
last_reboot.reason.code	string	query	False	Filter by last_reboot.reason.code

Name	Type	In	Required	Description
last_reboot.reason.target	string	query	False	Filter by last_reboot.reason.target • Introduced in: 9.10
name	string	query	False	Filter by name
firmware_version	string	query	False	Filter by firmware_version
dram_single_bit_error_count	integer	query	False	Filter by dram_single_bit_error_count
chassis_throughput_state	string	query	False	Filter by chassis_throughput_state
symbolic_name	string	query	False	Filter by symbolic_name
temperature_sensor.minimum	integer	query	False	Filter by temperature_sensor.minimum
temperature_sensor.maximum	integer	query	False	Filter by temperature_sensor.maximum
temperature_sensor.name	string	query	False	Filter by temperature_sensor.name
temperature_sensor.reading	integer	query	False	Filter by temperature_sensor.reading
temperature_sensor.state	string	query	False	Filter by temperature_sensor.state
sas_ports.id	integer	query	False	Filter by sas_ports.id

Name	Type	In	Required	Description
sas_ports.phy_1.state	string	query	False	Filter by sas_ports.phy_1.state
sas_ports.data_rate_capability	number	query	False	Filter by sas_ports.data_rate_capability
sas_ports.negotiated_data_rate	number	query	False	Filter by sas_ports.negotiated_data_rate
sas_ports.phy_3.state	string	query	False	Filter by sas_ports.phy_3.state
sas_ports.phy_2.state	string	query	False	Filter by sas_ports.phy_2.state
sas_ports.phy_4.state	string	query	False	Filter by sas_ports.phy_4.state
sas_ports.enabled	boolean	query	False	Filter by sas_ports.enabled
sas_ports.wwn	string	query	False	Filter by sas_ports.wwn
sas_ports.state	string	query	False	Filter by sas_ports.state
sas_ports.cable.vendor	string	query	False	Filter by sas_ports.cable.vendor
sas_ports.cable.serial_number	string	query	False	Filter by sas_ports.cable.serial_number
sas_ports.cable.part_number	string	query	False	Filter by sas_ports.cable.part_number

Name	Type	In	Required	Description
sas_ports.cable.technology	string	query	False	Filter by sas_ports.cable.technology
power_supply_units.name	string	query	False	Filter by power_supply_units.name
power_supply_units.state	string	query	False	Filter by power_supply_units.state
errors.component.id	integer	query	False	Filter by errors.component.id
errors.component.unique_id	string	query	False	Filter by errors.component.unique_id
errors.component.name	string	query	False	Filter by errors.component.name
errors.reason.arguments.message	string	query	False	Filter by errors.reason.arguments.message • Introduced in: 9.10
errors.reason.arguments.code	string	query	False	Filter by errors.reason.arguments.code • Introduced in: 9.10
errors.reason.message	string	query	False	Filter by errors.reason.message
errors.reason.code	string	query	False	Filter by errors.reason.code

Name	Type	In	Required	Description
errors.reason.target	string	query	False	Filter by errors.reason.target • Introduced in: 9.10
errors.severity	string	query	False	Filter by errors.severity
errors.type	string	query	False	Filter by errors.type
security_enabled	boolean	query	False	Filter by security_enabled
ip_address	string	query	False	Filter by ip_address
serial_number	string	query	False	Filter by serial_number
model	string	query	False	Filter by model
fc_ports.enabled	boolean	query	False	Filter by fc_ports.enabled
fc_ports.negotiated_data_rate	number	query	False	Filter by fc_ports.negotiated_data_rate
fc_ports.data_rate_capability	number	query	False	Filter by fc_ports.data_rate_capability
fc_ports.state	string	query	False	Filter by fc_ports.state
fc_ports.wwn	string	query	False	Filter by fc_ports.wwn
fc_ports.peer_wwn	string	query	False	Filter by fc_ports.peer_wwn
fc_ports.configured_data_rate	number	query	False	Filter by fc_ports.configured_data_rate
fc_ports.id	integer	query	False	Filter by fc_ports.id

Name	Type	In	Required	Description
fc_ports.connection_mode	string	query	False	Filter by fc_ports.connection_mode
fc_ports.sfp.vendor	string	query	False	Filter by fc_ports.sfp.vendor
fc_ports.sfp.data_rate_capability	number	query	False	Filter by fc_ports.sfp.data_rate_capability
fc_ports.sfp.part_number	string	query	False	Filter by fc_ports.sfp.part_number
fc_ports.sfp.serial_number	string	query	False	Filter by fc_ports.sfp.serial_number
monitoring_enabled	boolean	query	False	Filter by monitoring_enabled
vendor	string	query	False	Filter by vendor
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • 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[storage_bridge]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "chassis_throughput_state": "ok",
    "errors": {
      "reason": {
        "arguments": {
          "code": "string",
          "message": "string"
        },
        "code": "4",
        "message": "entry doesn't exist",
        "target": "uuid"
      },
      "severity": "unknown",
      "type": "unknown"
    },
    "fc_ports": {
      "configured_data_rate": 0,
      "connection_mode": "loop",
      "data_rate_capability": 2,
      "negotiated_data_rate": 0,
      "peer_wwn": "200650eb1a238892",
      "sfp": {
        "data_rate_capability": 2
      },
      "state": "error",
      "wwn": "2100001086a54100"
    },
    "firmware_version": "4.10 007A",
    "ip_address": "string",
    "last_reboot": {
      "reason": {
        "arguments": {
          "code": "string",
          "message": "string"
        }
      }
    }
  }
}
```

```

    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "time": "2020-12-09 05:47:58 +0000"
},
"managed_by": "snmp",
"model": "FibreBridge6500N",
"name": "ATTO_FibreBridge6500N_1",
"paths": {
  "name": "2c",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "source_port": {
    "id": "100050eb1a238892",
    "name": "rtp-fc03-41kk11:1"
  },
  "target_port": {
    "id": "100050eb1a238892",
    "name": "rtp-fc03-41kk11:6",
    "wwn": "2100001086a54100"
  }
},
"power_supply_units": {
  "state": "ok"
},
"sas_ports": {
  "data_rate_capability": 0,
  "negotiated_data_rate": 0,
  "state": "error",
  "wwn": "2100001086a54100"
},
"serial_number": "FB7600N100004",
"state": "unknown",
"symbolic_name": "rtp-fcsas03-41kk11",
"temperature_sensor": {
  "name": "Chassis temperature sensor",
  "state": "ok"
}

```



```
    },
    "vendor": "atto",
    "wwn": 2000001086600476
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

component

Name	Type	Description
id	integer	Bridge error component ID
name	string	Bridge error component name
unique_id	string	Bridge error component unique ID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

errors

Name	Type	Description
component	component	
reason	error	
severity	string	Bridge error severity
type	string	Bridge error type

sfp

Name	Type	Description
data_rate_capability	number	Bridge FC port SFP data rate capability, in Gbps
part_number	string	
serial_number	string	Bridge FC port SFP serial number
vendor	string	Bridge FC port SFP vendor

fc_ports

Name	Type	Description
configured_data_rate	number	Bridge FC port configured data rate, in Gbps
connection_mode	string	Bridge FC port configured connection mode
data_rate_capability	number	Bridge FC port data rate capability, in Gbps
enabled	boolean	Indicates whether the bridge FC port is enabled.
id	integer	Bridge FC port index
negotiated_data_rate	number	Bridge FC port negotiated data rate, in Gbps
peer_wwn	string	Bridge FC port peer port world wide name
sfp	sfp	

Name	Type	Description
state	string	Bridge FC port state
wwn	string	Bridge FC port world wide name

error

The error message and code explaining why the bridge rebooted.

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.

last_reboot

Name	Type	Description
reason	error	The error message and code explaining why the bridge rebooted.
time	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

source_port

Name	Type	Description
id	string	Initiator side switch port id
name	string	Initiator side switch port name

target_port

Name	Type	Description
id	string	Target side switch port id
name	string	Target side switch port name
wwn	string	Target side switch port world wide name

paths

Name	Type	Description
name	string	
node	node	
source_port	source_port	
target_port	target_port	

power_supply_units

Name	Type	Description
name	string	Power supply unit name
state	string	Power supply unit state

cable

Name	Type	Description
part_number	string	Bridge cable part number
serial_number	string	Bridge cable serial number
technology	string	Bridge cable type
vendor	string	Bridge cable vendor

phy_1

Name	Type	Description
state	string	Bridge SAS port PHY1 state

phy_2

Name	Type	Description
state	string	Bridge SAS port PHY2 state

phy_3

Name	Type	Description
state	string	Bridge SAS port PHY3 state

phy_4

Name	Type	Description
state	string	Bridge SAS port PHY4 state

sas_ports

Name	Type	Description
cable	cable	
data_rate_capability	number	Bridge SAS port data rate capability, in Gbps
enabled	boolean	Indicates whether a bridge SAS port is enabled.
id	integer	Bridge SAS port index
negotiated_data_rate	number	Bridge SAS port negotiated data rate, in Gbps
phy_1	phy_1	
phy_2	phy_2	
phy_3	phy_3	
phy_4	phy_4	
state	string	Bridge SAS port state

Name	Type	Description
wwn	string	Bridge SAS port world wide name

temperature_sensor

Name	Type	Description
maximum	integer	Maximum safe operating temperature, in degrees Celsius.
minimum	integer	Minimum safe operating temperature, in degrees Celsius.
name	string	Temperature sensor name
reading	integer	Chassis temperature sensor reading, in degrees Celsius.
state	string	

storage_bridge

Name	Type	Description
chassis_throughput_state	string	Chassis throughput status
dram_single_bit_error_count	integer	
errors	array[errors]	
fc_ports	array[fc_ports]	
firmware_version	string	Bridge firmware version
ip_address	string	IP Address
last_reboot	last_reboot	
managed_by	string	
model	string	Bridge model
monitoring_enabled	boolean	Indicates whether monitoring is enabled for the bridge.
name	string	Bridge name
paths	array[paths]	
power_supply_units	array[power_supply_units]	

Name	Type	Description
sas_ports	array[sas_ports]	
security_enabled	boolean	Indicates whether security is enabled for the bridge.
serial_number	string	Bridge serial number
state	string	Bridge state
symbolic_name	string	Bridge symbolic name
temperature_sensor	temperature_sensor	
vendor	string	Bridge vendor
wwn	string	Bridge world wide name

Retrieve a specific bridge

GET /storage/bridges/{wwn}

Introduced In: 9.9

Retrieves a specific bridge

Related ONTAP commands

- `storage bridge show`

Learn more

- [DOC /storage/bridges](#)

Parameters

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

Response

Status: 200, Ok

Name	Type	Description
chassis_throughput_state	string	Chassis throughput status
dram_single_bit_error_count	integer	
errors	array[errors]	
fc_ports	array[fc_ports]	
firmware_version	string	Bridge firmware version
ip_address	string	IP Address
last_reboot	last_reboot	
managed_by	string	
model	string	Bridge model
monitoring_enabled	boolean	Indicates whether monitoring is enabled for the bridge.
name	string	Bridge name
paths	array[paths]	
power_supply_units	array[power_supply_units]	
sas_ports	array[sas_ports]	
security_enabled	boolean	Indicates whether security is enabled for the bridge.
serial_number	string	Bridge serial number
state	string	Bridge state
symbolic_name	string	Bridge symbolic name
temperature_sensor	temperature_sensor	
vendor	string	Bridge vendor
wwn	string	Bridge world wide name

Example response

```
{
  "chassis_throughput_state": "ok",
  "errors": {
    "reason": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    },
    "severity": "unknown",
    "type": "unknown"
  },
  "fc_ports": {
    "configured_data_rate": 0,
    "connection_mode": "loop",
    "data_rate_capability": 2,
    "negotiated_data_rate": 0,
    "peer_wwn": "200650eb1a238892",
    "sfp": {
      "data_rate_capability": 2
    },
    "state": "error",
    "wwn": "2100001086a54100"
  },
  "firmware_version": "4.10 007A",
  "ip_address": "string",
  "last_reboot": {
    "reason": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    },
    "time": "2020-12-09 05:47:58 +0000"
  },
  "managed_by": "snmp",
  "model": "FibreBridge6500N",
  "name": "ATTO_FibreBridge6500N_1",
}
```

```

"paths": {
  "name": "2c",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "source_port": {
    "id": "100050eb1a238892",
    "name": "rtp-fc03-41kk11:1"
  },
  "target_port": {
    "id": "100050eb1a238892",
    "name": "rtp-fc03-41kk11:6",
    "wn": "2100001086a54100"
  }
},
"power_supply_units": {
  "state": "ok"
},
"sas_ports": {
  "data_rate_capability": 0,
  "negotiated_data_rate": 0,
  "state": "error",
  "wn": "2100001086a54100"
},
"serial_number": "FB7600N100004",
"state": "unknown",
"symbolic_name": "rtp-fcsas03-41kk11",
"temperature_sensor": {
  "name": "Chassis temperature sensor",
  "state": "ok"
},
"vendor": "atto",
"wn": 2000001086600476
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

component

Name	Type	Description
id	integer	Bridge error component ID
name	string	Bridge error component name
unique_id	string	Bridge error component unique ID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

errors

Name	Type	Description
component	component	
reason	error	
severity	string	Bridge error severity
type	string	Bridge error type

sfp

Name	Type	Description
data_rate_capability	number	Bridge FC port SFP data rate capability, in Gbps
part_number	string	
serial_number	string	Bridge FC port SFP serial number
vendor	string	Bridge FC port SFP vendor

fc_ports

Name	Type	Description
configured_data_rate	number	Bridge FC port configured data rate, in Gbps
connection_mode	string	Bridge FC port configured connection mode
data_rate_capability	number	Bridge FC port data rate capability, in Gbps
enabled	boolean	Indicates whether the bridge FC port is enabled.
id	integer	Bridge FC port index
negotiated_data_rate	number	Bridge FC port negotiated data rate, in Gbps
peer_wwn	string	Bridge FC port peer port world wide name
sfp	sfp	
state	string	Bridge FC port state
wwn	string	Bridge FC port world wide name

error

The error message and code explaining why the bridge rebooted.

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.

last_reboot

Name	Type	Description
reason	error	The error message and code explaining why the bridge rebooted.
time	string	

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

source_port

Name	Type	Description
id	string	Initiator side switch port id
name	string	Initiator side switch port name

target_port

Name	Type	Description
id	string	Target side switch port id
name	string	Target side switch port name
wwn	string	Target side switch port world wide name

paths

Name	Type	Description
name	string	
node	node	
source_port	source_port	
target_port	target_port	

power_supply_units

Name	Type	Description
name	string	Power supply unit name
state	string	Power supply unit state

cable

Name	Type	Description
part_number	string	Bridge cable part number
serial_number	string	Bridge cable serial number
technology	string	Bridge cable type
vendor	string	Bridge cable vendor

phy_1

Name	Type	Description
state	string	Bridge SAS port PHY1 state

phy_2

Name	Type	Description
state	string	Bridge SAS port PHY2 state

phy_3

Name	Type	Description
state	string	Bridge SAS port PHY3 state

phy_4

Name	Type	Description
state	string	Bridge SAS port PHY4 state

sas_ports

Name	Type	Description
cable	cable	
data_rate_capability	number	Bridge SAS port data rate capability, in Gbps
enabled	boolean	Indicates whether a bridge SAS port is enabled.
id	integer	Bridge SAS port index
negotiated_data_rate	number	Bridge SAS port negotiated data rate, in Gbps
phy_1	phy_1	
phy_2	phy_2	
phy_3	phy_3	
phy_4	phy_4	
state	string	Bridge SAS port state
wwn	string	Bridge SAS port world wide name

temperature_sensor

Name	Type	Description
maximum	integer	Maximum safe operating temperature, in degrees Celsius.
minimum	integer	Minimum safe operating temperature, in degrees Celsius.
name	string	Temperature sensor name
reading	integer	Chassis temperature sensor reading, in degrees Celsius.
state	string	

Report cluster-wide storage details across different tiers

GET /storage/cluster

Introduced In: 9.6

Reports cluster wide storage details across different tiers. By default, this endpoint returns all fields. Supports the following roles: admin, and readonly.

Parameters

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

Response

Status: 200, Ok

Name	Type	Description
block_storage	block_storage	
cloud_storage	cloud_storage	
efficiency	space_efficiency	Storage efficiency
efficiency_without_snapshots	space_efficiency	Storage efficiency that does not include the savings provided by Snapshot copies.

Name	Type	Description
efficiency_without_snapshots_flexclones	space_efficiency	Storage efficiency that does not include the savings provided by Snapshot copies and FlexClones.

Example response

```
{
  "block_storage": {
    "medias": {
      "efficiency": {
        "logical_used": 0,
        "ratio": 0,
        "savings": 0
      },
      "efficiency_without_snapshots": {
        "logical_used": 0,
        "ratio": 0,
        "savings": 0
      },
      "efficiency_without_snapshots_flexclones": {
        "logical_used": 0,
        "ratio": 0,
        "savings": 0
      },
      "type": "hdd"
    }
  },
  "cloud_storage": {
    "used": 0
  },
  "efficiency": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots_flexclones": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

space_efficiency

Storage Efficiency

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

space_efficiency

Storage efficiency that does not include the savings provided by Snapshot copies.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

space_efficiency

Storage efficiency that does not include the savings provided by Snapshot copies and FlexClones.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

medias

Name	Type	Description
available	integer	Available space

Name	Type	Description
efficiency	space_efficiency	Storage Efficiency
efficiency_without_snapshots	space_efficiency	Storage efficiency that does not include the savings provided by Snapshot copies.
efficiency_without_snapshots_flex_clones	space_efficiency	Storage efficiency that does not include the savings provided by Snapshot copies and FlexClones.
physical_used	integer	Total physical used space
size	integer	Total space
type	string	The type of media being used
used	integer	Used space

block_storage

Name	Type	Description
available	integer	Available space across the cluster
inactive_data	integer	Inactive data across all aggregates
medias	array[medias]	
physical_used	integer	Total physical used space across the cluster
size	integer	Total space across the cluster
used	integer	Space used (includes volume reserves)

cloud_storage

Name	Type	Description
used	integer	Total space used in cloud.

space_efficiency

Storage efficiency

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage disks

Storage disks endpoint overview

Retrieving storage disk information

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

Examples

1) Retrieve a list of disks from the cluster.

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

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

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

# The response:
{
  "records": [
    {
      "name": "1.24.4",
      "_links": {
        "self": {
          "href": "/api/storage/disks/1.24.4"
        }
      }
    },
    {
      "name": "1.24.3",
      "_links": {
        "self": {
          "href": "/api/storage/disks/1.24.3"
        }
      }
    },
    {
      "name": "1.24.5",
      "_links": {
        "self": {
          "href": "/api/storage/disks/1.24.5"
        }
      }
    },
    {
      "name": "1.24.0",
      "_links": {
        "self": {
          "href": "/api/storage/disks/1.24.0"
        }
      }
    }
  ]
}
```

```

    },
    {
      "name": "1.24.2",
      "_links": {
        "self": {
          "href": "/api/storage/disks/1.24.2"
        }
      }
    },
    {
      "name": "1.24.1",
      "_links": {
        "self": {
          "href": "/api/storage/disks/1.24.1"
        }
      }
    }
  ],
  "num_records": 6,
  "_links": {
    "self": {
      "href": "/api/storage/disks"
    }
  }
}

```

2) Retrieve a specific disk from the cluster.

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

```

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

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

# The response:
{
  "name": "1.24.3",
  "uid":
  "50000394:0808AA88:00000000:00000000:00000000:00000000:00000000:00000000:0

```

```

0000000:00000000",
"serial_number": "EC47PC5021SW",
"model": "X421_FAL12450A10",
"vendor": "NETAPP",
"firmware_version": "NA02",
"usable_size": 438304768000,
"rpm": 10000,
"type": "sas",
"effective_type": "sas",
"class": "performance",
"container_type": "aggregate",
"pool": "pool0",
"state": "present",
"node": {
  "uuid": "3a89ed49-8c6d-11e8-93bc-00a0985a64b6",
  "name": "node-2",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/3a89ed49-8c6d-11e8-93bc-00a0985a64b6"
    }
  }
},
"home_node": {
  "uuid": "3a89ed49-8c6d-11e8-93bc-00a0985a64b6",
  "name": "node-2",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/3a89ed49-8c6d-11e8-93bc-00a0985a64b6"
    }
  }
},
"aggregates": [
  {
    "uuid": "3fd9c345-ba91-4949-a7b1-6e2b898d74e3",
    "name": "node_2_SAS_1",
    "_links": {
      "self": {
        "href": "/api/storage/aggregates/3fd9c345-ba91-4949-a7b1-
6e2b898d74e3"
      }
    }
  }
],
"shelf": {
  "uid": "10318311901725526608",
  "_links": {

```

```

    "self": {
      "href": "/api/storage/shelves/10318311901725526608"
    }
  },
  "local": true,
  "paths": [
    {
      "initiator": "3a",
      "port_name": "B",
      "port_type": "sas",
      "wwnn": "5000cca02f0e6768",
      "wwpn": "5000cca02f0e676a",
      "node": {
        "name": "vsim3",
        "uuid": "3f7fa09a-5c56-11ec-b366-005056bbbb3f",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/3f7fa09a-5c56-11ec-b366-
005056bbbb3f"
          }
        }
      }
    },
    {
      "initiator": "3d",
      "port_name": "A",
      "port_type": "sas",
      "wwnn": "5000cca02f0e6768",
      "wwpn": "5000cca02f0e6769",
      "node": {
        "name": "vsim4",
        "uuid": "4f7fa09a-5c56-11ec-b366-005056bbbb3f",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/4f7fa09a-5c56-11ec-b366-
005056bbbb3f"
          }
        }
      }
    },
    {
      "initiator": "3d",
      "port_name": "A",
      "port_type": "sas",
      "wwnn": "5000cca02f0e6768",

```

```

"wwpn": "5000cca02f0e6769",
"node": {
  "name": "vsim3",
  "uuid": "3f7fa09a-5c56-11ec-b366-005056bbbb3f",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/3f7fa09a-5c56-11ec-b366-
005056bbbb3f"
    }
  }
},
{
  "initiator": "3a",
  "port_name": "B",
  "port_type": "sas",
  "wwnn": "5000cca02f0e6768",
  "wwpn": "5000cca02f0e676a",
  "node": {
    "name": "vsim4",
    "uuid": "4f7fa09a-5c56-11ec-b366-005056bbbb3f",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/4f7fa09a-5c56-11ec-b366-
005056bbbb3f"
      }
    }
  }
},
"outage": {
  "persistently_failed": true,
  "reason": {
    "message": "Failed disk. Reason: \"admin failed\".",
    "code": "721081"
  }
},
"bay": 3,
"_links": {
  "self": {
    "href": "/api/storage/disks/1.24.3"
  }
},
"error": [
{
  "reason": {

```

```

    "message": "\"The node is configured with All-Flash Optimized
personality and this disk is not an SSD. The disk needs to be removed from
the system.\",
    \"code\": \"721082\"
  },
  \"type\": \"notallflashdisk\"
}
],
\"bytes_per_sector\": 512,
\"sector_count\": 1172123568,
\"right_size_sector_count\": 5579776,
\"physical_size\": 438804988000,
\"stats\": {
  \"average_latency\": 6,
  \"throughput\": 1957888,
  \"iops_total\": 12854,
  \"path_error_count\": 0,
  \"power_on_hours\": 11797
}
}

```

3) Retrieving a specific disk from the hypervisor

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

```

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

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/disks/NET-3.2" -H "accept:
application/hal+json"

# The response:
{
  \"name\": \"NET-3.2\",
  \"uid\":
\"32343637:65386464:00000000:00000000:00000000:00000000:00000000:00000000:0
0000000:00000000\",
  \"serial_number\": \"3234363765386464\",
  \"model\": \"PHA-DISK\",
  \"vendor\": \"NETAPP\",
  \"firmware_version\": \"0001\",

```

```

"type": "vmdisk",
"class": "virtual",
"container_type": "mediator",
"pool": "pool0",
"node": {
  "uuid": "f4cb78ba-5841-11ec-80c4-916f62b4cd44",
  "name": "aws-43905099-44129379-awsha-vm1",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/f4cb78ba-5841-11ec-80c4-
916f62b4cd44"
    }
  }
},
"home_node": {
  "uuid": "f4cb78ba-5841-11ec-80c4-916f62b4cd44",
  "name": "aws-43905099-44129379-awsha-vm1",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/f4cb78ba-5841-11ec-80c4-
916f62b4cd44"
    }
  }
},
"local": true,
"paths": [
{
  "initiator": "0f",
  "port_name": "A",
  "port_type": "sas",
  "wwnn": "53059d50444f5476",
  "wwpn": "53059d50444f5476",
  "vmdisk_hypervisor_file_name": "LUN 4.0",
  "node": {
    "name": "aws-43905099-44129379-awsha-vm1",
    "uuid": "f4cb78ba-5841-11ec-80c4-916f62b4cd44",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/f4cb78ba-5841-11ec-80c4-
916f62b4cd44"
      }
    }
  }
},
{
  "initiator": "0f",

```

```

    "port_name": "A",
    "port_type": "sas",
    "wwnn": "53059d50444f5476",
    "wwpn": "53059d50444f5476",
    "vmdisk_hypervisor_file_name": "LUN 2.0",
    "node": {
      "name": "aws-43905099-44129379-awsha-vm1",
      "uuid": "f4cb78ba-5841-11ec-80c4-916f62b4cd44",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/f4cb78ba-5841-11ec-80c4-
916f62b4cd44"
        }
      }
    }
  ],
  "outage": {
    "persistently_failed": false,
    "reason": {
      "message": "Failed disk. Reason: \"\".",
      "code": "721081"
    }
  },
  "self_encrypting": false,
  "fips_certified": false,
  "bytes_per_sector": 512,
  "sector_count": 204808,
  "right_size_sector_count": 5579776,
  "physical_size": 204808,
  "stats": {
    "average_latency": 2157188883,
    "throughput": 4096,
    "iops_total": 1,
    "path_error_count": 0,
    "power_on_hours": 0
  },
  "_links": {
    "self": {
      "href": "/api/storage/disks/NET-3.2"
    }
  }
}

```


Modifying storage disk

The storage disk PATCH API modifies disk ownership, unfails a disk, or updates encrypting drive authentication keys (AKs) in the cluster. The storage disk API currently supports patching one attribute at a time.

Updating the disk ownership for a specified disk

1. When the disk is not assigned

When the disk is a spare (or unowned) disk and node name is specified, the PATCH operation assigns the disk to the specified node. Optionally, pool name can also be specified along with node name. Accepted pool names are: pool0, pool1.

2. When the disk is already assigned

When the disk is already assigned (already has a owner), and a new node is specified, the PATCH operation changes the ownership to the new node. Optionally, pool name can also be specified along with node name. Accepted pool names are: pool0, pool1.

Removing the disk ownership for a specified disk

When the disk is already assigned, and node name is specified as null (no-quotes), the PATCH operation removes the owner.

Examples

1. Update the disk ownership for an unowned disk

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-name>" -H
"accept: application/hal+json" -H "Content-Type: application/hal+json" -d
'{"node": {"name": "node-name"}}'

# The response:
{
}
```

2. Update the disk ownership for an already owned disk

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-name>" -H
"accept: application/hal+json" -H "Content-Type: application/hal+json" -d
'{"node": {"name": "node-name"}}'

# The response:
{
}
```

3. Update the disk pool for a disk (can be either owned or unowned).

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-name>" -H
"accept: application/hal+json" -H "Content-Type: application/hal+json" -d
'{"node": {"name": "node-name"}, "pool": "pool0"}'

# The response:
{
}
```

4. Rekey the data authentication key (AK) of all encrypting drives to an authentication key (AK) selected automatically by the system

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

# The call:
curl -X PATCH "https://<mgmt-
ip>/api/storage/disks?name=*&encryption_operation=rekey_data_auto_id" -H
"accept: application/hal+json" -H "Content-Type: application/hal+json"

# The response contains the number of disks attempted.
{
  "num_records": 32
}
```

5. Cryptographically sanitize a spare or broken disk

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-
name>&encryption_operation=sanitize_disk" -H "accept:
application/hal+json" -H "Content-Type: application/hal+json"

# The response contains the number of disks attempted.
{
  "num_records": 1
}
```

6. Unfailing a disk to a spare.

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-name>" -d
'{"state": "spare"}' -H "accept: application/hal+json" -H "Content-Type:
application/hal+json"

# The response:
{
}
```

7. Unfailing a disk and attempting to reassimilate filesystem labels.

If unable or unnecessary to reassimilate filesystem labels, the disk will be set as spare.

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/disks?name=<disk-name>" -d
'{"state": "present"}' -H "accept: application/hal+json" -H "Content-Type:
application/hal+json"

# The response:
{
}
```

Retrieve a collection of disks

GET /storage/disks

Introduced In: 9.6

Retrieves a collection of disks.

Related ONTAP commands

- storage disk show

Learn more

- [DOC /storage/disks](#)

Parameters

Name	Type	In	Required	Description
protection_mode	string	query	False	Filter by protection_mode <ul style="list-style-type: none">• Introduced in: 9.7
self_encrypting	boolean	query	False	Filter by self_encrypting <ul style="list-style-type: none">• Introduced in: 9.7
state	string	query	False	Filter by state
name	string	query	False	Filter by name
firmware_version	string	query	False	Filter by firmware_version
bytes_per_sector	integer	query	False	Filter by bytes_per_sector <ul style="list-style-type: none">• Introduced in: 9.9
vendor	string	query	False	Filter by vendor
home_node.uuid	string	query	False	Filter by home_node.uuid
home_node.name	string	query	False	Filter by home_node.name
container_type	string	query	False	Filter by container_type
pool	string	query	False	Filter by pool

Name	Type	In	Required	Description
outage.reason.arguments.message	string	query	False	Filter by outage.reason.arguments.message • Introduced in: 9.10
outage.reason.arguments.code	string	query	False	Filter by outage.reason.arguments.code • Introduced in: 9.10
outage.reason.message	string	query	False	Filter by outage.reason.message • Introduced in: 9.9
outage.reason.code	string	query	False	Filter by outage.reason.code • Introduced in: 9.9
outage.reason.target	string	query	False	Filter by outage.reason.target • Introduced in: 9.10
outage.persistently_failed	boolean	query	False	Filter by outage.persistently_failed • Introduced in: 9.9
error.type	string	query	False	Filter by error.type • Introduced in: 9.9

Name	Type	In	Required	Description
error.reason.arguments.message	string	query	False	Filter by error.reason.arguments.message • Introduced in: 9.10
error.reason.arguments.code	string	query	False	Filter by error.reason.arguments.code • Introduced in: 9.10
error.reason.message	string	query	False	Filter by error.reason.message • Introduced in: 9.9
error.reason.code	string	query	False	Filter by error.reason.code • Introduced in: 9.9
error.reason.target	string	query	False	Filter by error.reason.target • Introduced in: 9.10
storage_pool.uuid	string	query	False	Filter by storage_pool.uuid • Introduced in: 9.11
storage_pool.name	string	query	False	Filter by storage_pool.name • Introduced in: 9.11

Name	Type	In	Required	Description
right_size_sector_count	integer	query	False	Filter by right_size_sector_count • Introduced in: 9.11
drawer.slot	integer	query	False	Filter by drawer.slot
drawer.id	integer	query	False	Filter by drawer.id
dr_node.name	string	query	False	Filter by dr_node.name
dr_node.uuid	string	query	False	Filter by dr_node.uuid
effective_type	string	query	False	Filter by effective_type • Introduced in: 9.9
local	boolean	query	False	Filter by local • Introduced in: 9.9
overall_security	string	query	False	Filter by overall_security • Introduced in: 9.11
type	string	query	False	Filter by type
stats.throughput	integer	query	False	Filter by stats.throughput • Introduced in: 9.9
stats.power_on_hours	integer	query	False	Filter by stats.power_on_hours • Introduced in: 9.9

Name	Type	In	Required	Description
stats.path_error_count	integer	query	False	Filter by stats.path_error_count • Introduced in: 9.9
stats.average_latency	integer	query	False	Filter by stats.average_latency • Introduced in: 9.9
stats.iops_total	integer	query	False	Filter by stats.iops_total • Introduced in: 9.9
paths.port_name	string	query	False	Filter by paths.port_name • Introduced in: 9.9
paths.wwnn	string	query	False	Filter by paths.wwnn • Introduced in: 9.9
paths.port_type	string	query	False	Filter by paths.port_type • Introduced in: 9.9
paths.initiator	string	query	False	Filter by paths.initiator • Introduced in: 9.9
paths.vmdisk_hypervisor_file_name	string	query	False	Filter by paths.vmdisk_hypervisor_file_name • Introduced in: 9.11

Name	Type	In	Required	Description
paths.node.uuid	string	query	False	Filter by paths.node.uuid • Introduced in: 9.11
paths.node.name	string	query	False	Filter by paths.node.name • Introduced in: 9.11
paths.wwpn	string	query	False	Filter by paths.wwpn • Introduced in: 9.9
bay	integer	query	False	Filter by bay
class	string	query	False	Filter by class
rated_life_used_percent	integer	query	False	Filter by rated_life_used_percent
sector_count	integer	query	False	Filter by sector_count • Introduced in: 9.9
shelf.uid	string	query	False	Filter by shelf.uid
compliance_standard	string	query	False	Filter by compliance_standard • Introduced in: 9.11
physical_size	integer	query	False	Filter by physical_size • Introduced in: 9.11
node.uuid	string	query	False	Filter by node.uuid

Name	Type	In	Required	Description
node.name	string	query	False	Filter by node.name
virtual.storage_account	string	query	False	Filter by virtual.storage_account • Introduced in: 9.11
virtual.target_addresses	string	query	False	Filter by virtual.target_addresses • Introduced in: 9.13
virtual.object	string	query	False	Filter by virtual.object • Introduced in: 9.11
virtual.container	string	query	False	Filter by virtual.container • Introduced in: 9.11
rpm	integer	query	False	Filter by rpm
uid	string	query	False	Filter by uid
control_standard	string	query	False	Filter by control_standard • Introduced in: 9.11
aggregates.name	string	query	False	Filter by aggregates.name
aggregates.uuid	string	query	False	Filter by aggregates.uuid
usable_size	integer	query	False	Filter by usable_size

Name	Type	In	Required	Description
serial_number	string	query	False	Filter by serial_number
fips_certified	boolean	query	False	Filter by fips_certified • Introduced in: 9.7
model	string	query	False	Filter by model
key_id.data	string	query	False	Filter by key_id.data • Introduced in: 9.7
key_id.fips	string	query	False	Filter by key_id.fips • Introduced in: 9.7
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • 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[disk]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "aggregates": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "bay": 1,
    "bytes_per_sector": 520,
    "class": "solid_state",
    "compliance_standard": "FIPS 140-2",
    "container_type": "spare",
    "control_standard": "TCG Enterprise",
    "dr_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "effective_type": "vmdisk",
    "error": {
      "reason": {
        "arguments": {
          "code": "string",
          "message": "string"
        },
        "code": "4",
        "message": "entry doesn't exist",
        "target": "uuid"
      },
      "type": "notallflashdisk"
    },
    "firmware_version": "NA51",
  }
}
```

```

"home_node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"model": "X421_HCOBE450A10",
"name": "1.0.1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"outage": {
  "reason": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
},
"overall_security": "Level 2",
"paths": {
  "initiator": "3a",
  "node.name": "vsim4",
  "node.uuid": "cf7fe057-526d-11ec-af4e-0050568e9df0",
  "port_name": "A",
  "port_type": "sas",
  "vmdisk_hypervisor_file_name": "xvds vol10a0567ae156ca59f6",
  "wwnn": "5000c2971c1b2b8c",
  "wwpn": "5000c2971c1b2b8d"
},
"physical_size": 228930,
"pool": "pool0",
"protection_mode": "data",
"rated_life_used_percent": 10,

```

```

"right_size_sector_count": 1172123568,
"rpm": 15000,
"sector_count": 1172123568,
"serial_number": "KHG2VX8R",
"shelf": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "uid": "7777841915827391056"
},
"state": "present",
"stats": {
  "average_latency": 3,
  "iops_total": 12854,
  "path_error_count": 0,
  "power_on_hours": 21016,
  "throughput": 1957888
},
"storage_pool": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "storage_pool_1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"type": "ssd",
"uid":
"002538E5:71B00B2F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000",
"usable_size": 959934889984,
"vendor": "NETAPP",
"virtual": {
  "container": "nviet12122018113936-rg",
  "object": "f1fu63se",
  "storage_account": "nviet12122018113936ps",
  "target_address": "string"
}
}
}

```


Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

aggregates

Name	Type	Description
_links	_links	
name	string	
uuid	string	

dr_node

Name	Type	Description
name	string	
uuid	string	

drawer

Name	Type	Description
id	integer	
slot	integer	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

The message and code detailing the error state of this disk.

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.

disk_error_info

Name	Type	Description
reason	error	The message and code detailing the error state of this disk.
type	string	Disk error type.

home_node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

key_id

Name	Type	Description
data	string	Key ID of the data authentication key
fips	string	Key ID of the FIPS authentication key

node

Name	Type	Description
_links	_links	
name	string	

Name	Type	Description
uuid	string	

error

This error message and code explaining the disk failure.

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.

outage

Indicates if a disk has an entry in the failed disk registry, along with the reason for the failure.

Name	Type	Description
persistently_failed	boolean	Indicates whether RAID maintains the state of this disk as failed across reboots.
reason	error	This error message and code explaining the disk failure.

disk_path_info

Name	Type	Description
initiator	string	Initiator port.
node.name	string	Controller with the initiator port for this path.
node.uuid	string	Controller UUID, to identify node for this path.
port_name	string	Name of the disk port.
port_type	string	Disk port type.

Name	Type	Description
vmdisk_hypervisor_file_name	string	Virtual disk hypervisor file name.
wwnn	string	Target device's World Wide Node Name.
wwpn	string	Target device's World Wide Port Name.

shelf

Shelf

Name	Type	Description
_links	_links	
uid	string	

stats

Name	Type	Description
average_latency	integer	Average I/O latency across all active paths, in milliseconds.
iops_total	integer	Total I/O operations per second read and written to this disk across all active paths.
path_error_count	integer	Disk path error count; failed I/O operations.
power_on_hours	integer	Hours powered on.
throughput	integer	Total disk throughput per second across all active paths, in bytes.

storage_pool

Shared Storage Pool

Name	Type	Description
_links	_links	
name	string	
uuid	string	

virtual

Information about backing storage for disks on cloud platforms.

Name	Type	Description
container	string	Container name of the virtual disk.
object	string	Object name of the virtual disk.
storage_account	string	Storage account name of the virtual disk.
target_address	string	Target address of the virtual disk.

disk

Name	Type	Description
aggregates	array[aggregates]	List of aggregates sharing this disk
bay	integer	Disk shelf bay
bytes_per_sector	integer	Bytes per sector.
class	string	Disk class
compliance_standard	string	Security standard that the device is certified to.
container_type	string	Type of overlying disk container
control_standard	string	Standard that the device supports for encryption control.
dr_node	dr_node	
drawer	drawer	
effective_type	string	Effective Disk type
encryption_operation	string	This field should only be set as a query parameter in a PATCH operation. It is input only and won't be returned by a subsequent GET.
error	array[disk_error_info]	List of disk errors information.

Name	Type	Description
fips_certified	boolean	
firmware_version	string	
home_node	home_node	
key_id	key_id	
local	boolean	Indicates if a disk is locally attached versus being remotely attached. A locally attached disk resides in the same proximity as the host cluster versus been attached to the remote cluster.
model	string	
name	string	Cluster-wide disk name
node	node	
outage	outage	Indicates if a disk has an entry in the failed disk registry, along with the reason for the failure.
overall_security	string	Overall Security rating, for FIPS-certified devices.
paths	array[disk_path_info]	List of paths to a disk
physical_size	integer	Physical size, in units of bytes
pool	string	Pool to which disk is assigned
protection_mode	string	Mode of drive data protection and FIPS compliance. Possible values are: <ul style="list-style-type: none"> • <i>open</i> - Data is unprotected • <i>data</i> - Data protection only, without FIPS compliance • <i>part</i> - Data is unprotected; other FIPS compliance settings present • <i>full</i> - Full data and FIPS compliance protection • <i>miss</i> - Protection mode information is not available

Name	Type	Description
rated_life_used_percent	integer	Percentage of rated life used
right_size_sector_count	integer	Number of usable disk sectors that remain after subtracting the right-size adjustment for this disk.
rpm	integer	Revolutions per minute
sector_count	integer	Number of sectors on the disk.
self_encrypting	boolean	
serial_number	string	
shelf	shelf	Shelf
state	string	State
stats	stats	
storage_pool	storage_pool	Shared Storage Pool
type	string	Disk interface type
uid	string	The unique identifier for a disk
usable_size	integer	
vendor	string	
virtual	virtual	Information about backing storage for disks on cloud platforms.

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 disk ownership, change authentication keys, or sanitize disks

PATCH /storage/disks

Introduced In: 9.7

Updates disk ownership, changes authentication keys, or sanitizes disks.

Related ONTAP commands

- `storage disk assign`
- `storage disk removeowner`
- `storage encryption disk modify -data-key-id`
- `storage encryption disk sanitize`
- `security key-manager key query -key-type NSE-AK`
- `storage disk unfailed`

Learn more

- [DOC /storage/disks](#)

Parameters

Name	Type	In	Required	Description
name	string	query	False	Disk name
node	string	query	False	Node to assign disk <ul style="list-style-type: none">• Introduced in: 9.8
pool	string	query	False	Pool to assign disk to <ul style="list-style-type: none">• Introduced in: 9.11

Name	Type	In	Required	Description
encryption_operation	string	query	False	<p>Name of the operation to apply to encrypting disks.</p> <ul style="list-style-type: none"> • rekey_data_default changes the data authentication key (AK) to the drive-unique Manufacture Secure ID (MSID) value. Allows the drive to be attached to other clusters. Disables data-at-rest protection without erasing the data. • rekey_data_auto_id changes the data authentication key (AK) to an AK the cluster selects automatically. Enables data-at-rest protection. • sanitize_disk cryptographically erases all user data from a spare or broken drive by altering the data encryption key. Resets the data AK to the drive-unique MSID value and disables data-at-rest protection. Used when a drive is being repurposed or returned. • enum: ["rekey_data_default", "rekey_data_auto_id",

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
aggregates	array[aggregates]	List of aggregates sharing this disk
bay	integer	Disk shelf bay
bytes_per_sector	integer	Bytes per sector.
class	string	Disk class
compliance_standard	string	Security standard that the device is certified to.
container_type	string	Type of overlying disk container
control_standard	string	Standard that the device supports for encryption control.
dr_node	dr_node	
drawer	drawer	
effective_type	string	Effective Disk type
encryption_operation	string	This field should only be set as a query parameter in a PATCH operation. It is input only and won't be returned by a subsequent GET.
error	array[disk_error_info]	List of disk errors information.
fips_certified	boolean	
firmware_version	string	
home_node	home_node	
key_id	key_id	

Name	Type	Description
local	boolean	Indicates if a disk is locally attached versus being remotely attached. A locally attached disk resides in the same proximity as the host cluster versus been attached to the remote cluster.
model	string	
name	string	Cluster-wide disk name
node	node	
outage	outage	Indicates if a disk has an entry in the failed disk registry, along with the reason for the failure.
overall_security	string	Overall Security rating, for FIPS-certified devices.
paths	array[disk_path_info]	List of paths to a disk
physical_size	integer	Physical size, in units of bytes
pool	string	Pool to which disk is assigned
protection_mode	string	Mode of drive data protection and FIPS compliance. Possible values are: <ul style="list-style-type: none"> • <i>open</i> - Data is unprotected • <i>data</i> - Data protection only, without FIPS compliance • <i>part</i> - Data is unprotected; other FIPS compliance settings present • <i>full</i> - Full data and FIPS compliance protection • <i>miss</i> - Protection mode information is not available
rated_life_used_percent	integer	Percentage of rated life used
right_size_sector_count	integer	Number of usable disk sectors that remain after subtracting the right-size adjustment for this disk.

Name	Type	Description
rpm	integer	Revolutions per minute
sector_count	integer	Number of sectors on the disk.
self_encrypting	boolean	
serial_number	string	
shelf	shelf	Shelf
state	string	State
stats	stats	
storage_pool	storage_pool	Shared Storage Pool
type	string	Disk interface type
uid	string	The unique identifier for a disk
usable_size	integer	
vendor	string	
virtual	virtual	Information about backing storage for disks on cloud platforms.

Example request

```
{
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resource/1"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "bay": 1,
  "bytes_per_sector": 520,
  "class": "solid_state",
  "compliance_standard": "FIPS 140-2",
  "container_type": "spare",
  "control_standard": "TCG Enterprise",
  "dr_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "effective_type": "vmdisk",
  "error": {
    "reason": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    },
    "type": "notallflashdisk"
  },
  "firmware_version": "NA51",
  "home_node": {
    "_links": {
      "self": {
        "href": "/api/resource/1"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "model": "X421_HCOBE450A10",
}
```

```

"name": "1.0.1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"outage": {
  "reason": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
},
"overall_security": "Level 2",
"paths": {
  "initiator": "3a",
  "node.name": "vsim4",
  "node.uuid": "cf7fe057-526d-11ec-af4e-0050568e9df0",
  "port_name": "A",
  "port_type": "sas",
  "vmdisk_hypervisor_file_name": "xvds vol10a0567ae156ca59f6",
  "wwnn": "5000c2971c1b2b8c",
  "wwpn": "5000c2971c1b2b8d"
},
"physical_size": 228930,
"pool": "pool0",
"protection_mode": "data",
"rated_life_used_percent": 10,
"right_size_sector_count": 1172123568,
"rpm": 15000,
"sector_count": 1172123568,
"serial_number": "KHG2VX8R",
"shelf": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},

```

```
    "uid": "7777841915827391056"
  },
  "state": "present",
  "stats": {
    "average_latency": 3,
    "iops_total": 12854,
    "path_error_count": 0,
    "power_on_hours": 21016,
    "throughput": 1957888
  },
  "storage_pool": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "storage_pool_1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "type": "ssd",
  "uid":
  "002538E5:71B00B2F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000",
  "usable_size": 959934889984,
  "vendor": "NETAPP",
  "virtual": {
    "container": "nviet12122018113936-rg",
    "object": "f1fu63se",
    "storage_account": "nviet12122018113936ps",
    "target_address": "string"
  }
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
720951	Unable to unfail the disk.
721066	Node is outside the list of controllers for disk.
721084	An invalid pool name was specified.
721085	The node name was not specified.
1441795	Setting the data key ID to the manufacture secure ID is not allowed when in FIPS-compliance mode.
14155777	The operation failed on one or more disks.
14155778	No self-encrypting disks were specified.
14155779	Status from a node shows that a conflicting operation has occurred. Some disk controls might have changed.
14155780	Could not retrieve the required key ID from the key manager.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregates

Name	Type	Description
_links	_links	
name	string	
uuid	string	

dr_node

Name	Type	Description
name	string	
uuid	string	

drawer

Name	Type	Description
id	integer	
slot	integer	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

The message and code detailing the error state of this disk.

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.

disk_error_info

Name	Type	Description
reason	error	The message and code detailing the error state of this disk.
type	string	Disk error type.

home_node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

key_id

Name	Type	Description
data	string	Key ID of the data authentication key
fips	string	Key ID of the FIPS authentication key

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

error

This error message and code explaining the disk failure.

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.

outage

Indicates if a disk has an entry in the failed disk registry, along with the reason for the failure.

Name	Type	Description
persistently_failed	boolean	Indicates whether RAID maintains the state of this disk as failed across reboots.
reason	error	This error message and code explaining the disk failure.

disk_path_info

Name	Type	Description
initiator	string	Initiator port.
node.name	string	Controller with the initiator port for this path.
node.uuid	string	Controller UUID, to identify node for this path.
port_name	string	Name of the disk port.
port_type	string	Disk port type.
vmdisk_hypervisor_file_name	string	Virtual disk hypervisor file name.
wwnn	string	Target device's World Wide Node Name.

Name	Type	Description
wwpn	string	Target device's World Wide Port Name.

shelf

Shelf

Name	Type	Description
_links	_links	
uid	string	

stats

Name	Type	Description
average_latency	integer	Average I/O latency across all active paths, in milliseconds.
iops_total	integer	Total I/O operations per second read and written to this disk across all active paths.
path_error_count	integer	Disk path error count; failed I/O operations.
power_on_hours	integer	Hours powered on.
throughput	integer	Total disk throughput per second across all active paths, in bytes.

storage_pool

Shared Storage Pool

Name	Type	Description
_links	_links	
name	string	
uuid	string	

virtual

Information about backing storage for disks on cloud platforms.

Name	Type	Description
container	string	Container name of the virtual disk.
object	string	Object name of the virtual disk.
storage_account	string	Storage account name of the virtual disk.
target_address	string	Target address of the virtual disk.

disk

Name	Type	Description
aggregates	array[aggregates]	List of aggregates sharing this disk
bay	integer	Disk shelf bay
bytes_per_sector	integer	Bytes per sector.
class	string	Disk class
compliance_standard	string	Security standard that the device is certified to.
container_type	string	Type of overlying disk container
control_standard	string	Standard that the device supports for encryption control.
dr_node	dr_node	
drawer	drawer	
effective_type	string	Effective Disk type
encryption_operation	string	This field should only be set as a query parameter in a PATCH operation. It is input only and won't be returned by a subsequent GET.
error	array[disk_error_info]	List of disk errors information.
fips_certified	boolean	

Name	Type	Description
firmware_version	string	
home_node	home_node	
key_id	key_id	
local	boolean	Indicates if a disk is locally attached versus being remotely attached. A locally attached disk resides in the same proximity as the host cluster versus been attached to the remote cluster.
model	string	
name	string	Cluster-wide disk name
node	node	
outage	outage	Indicates if a disk has an entry in the failed disk registry, along with the reason for the failure.
overall_security	string	Overall Security rating, for FIPS-certified devices.
paths	array[disk_path_info]	List of paths to a disk
physical_size	integer	Physical size, in units of bytes
pool	string	Pool to which disk is assigned
protection_mode	string	Mode of drive data protection and FIPS compliance. Possible values are: <ul style="list-style-type: none"> • <i>open</i> - Data is unprotected • <i>data</i> - Data protection only, without FIPS compliance • <i>part</i> - Data is unprotected; other FIPS compliance settings present • <i>full</i> - Full data and FIPS compliance protection • <i>miss</i> - Protection mode information is not available
rated_life_used_percent	integer	Percentage of rated life used

Name	Type	Description
right_size_sector_count	integer	Number of usable disk sectors that remain after subtracting the right-size adjustment for this disk.
rpm	integer	Revolutions per minute
sector_count	integer	Number of sectors on the disk.
self_encrypting	boolean	
serial_number	string	
shelf	shelf	Shelf
state	string	State
stats	stats	
storage_pool	storage_pool	Shared Storage Pool
type	string	Disk interface type
uid	string	The unique identifier for a disk
usable_size	integer	
vendor	string	
virtual	virtual	Information about backing storage for disks on cloud platforms.

Retrieve a specific disk

GET /storage/disks/{name}

Introduced In: 9.6

Retrieves a specific disk.

Related ONTAP commands

- `storage disk show`

Learn more

- [DOC /storage/disks](#)

Parameters

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

Response

Status: 200, Ok

Name	Type	Description
aggregates	array[aggregates]	List of aggregates sharing this disk
bay	integer	Disk shelf bay
bytes_per_sector	integer	Bytes per sector.
class	string	Disk class
compliance_standard	string	Security standard that the device is certified to.
container_type	string	Type of overlying disk container
control_standard	string	Standard that the device supports for encryption control.
dr_node	dr_node	
drawer	drawer	
effective_type	string	Effective Disk type
encryption_operation	string	This field should only be set as a query parameter in a PATCH operation. It is input only and won't be returned by a subsequent GET.
error	array[disk_error_info]	List of disk errors information.
fips_certified	boolean	
firmware_version	string	
home_node	home_node	

Name	Type	Description
key_id	key_id	
local	boolean	Indicates if a disk is locally attached versus being remotely attached. A locally attached disk resides in the same proximity as the host cluster versus been attached to the remote cluster.
model	string	
name	string	Cluster-wide disk name
node	node	
outage	outage	Indicates if a disk has an entry in the failed disk registry, along with the reason for the failure.
overall_security	string	Overall Security rating, for FIPS-certified devices.
paths	array[disk_path_info]	List of paths to a disk
physical_size	integer	Physical size, in units of bytes
pool	string	Pool to which disk is assigned
protection_mode	string	Mode of drive data protection and FIPS compliance. Possible values are: <ul style="list-style-type: none"> • <i>open</i> - Data is unprotected • <i>data</i> - Data protection only, without FIPS compliance • <i>part</i> - Data is unprotected; other FIPS compliance settings present • <i>full</i> - Full data and FIPS compliance protection • <i>miss</i> - Protection mode information is not available
rated_life_used_percent	integer	Percentage of rated life used

Name	Type	Description
right_size_sector_count	integer	Number of usable disk sectors that remain after subtracting the right-size adjustment for this disk.
rpm	integer	Revolutions per minute
sector_count	integer	Number of sectors on the disk.
self_encrypting	boolean	
serial_number	string	
shelf	shelf	Shelf
state	string	State
stats	stats	
storage_pool	storage_pool	Shared Storage Pool
type	string	Disk interface type
uid	string	The unique identifier for a disk
usable_size	integer	
vendor	string	
virtual	virtual	Information about backing storage for disks on cloud platforms.

Example response

```
{
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resource/1"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "bay": 1,
  "bytes_per_sector": 520,
  "class": "solid_state",
  "compliance_standard": "FIPS 140-2",
  "container_type": "spare",
  "control_standard": "TCG Enterprise",
  "dr_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "effective_type": "vmdisk",
  "error": {
    "reason": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    },
    "type": "notallflashdisk"
  },
  "firmware_version": "NA51",
  "home_node": {
    "_links": {
      "self": {
        "href": "/api/resource/1"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "model": "X421_HCOBE450A10",
}
```

```

"name": "1.0.1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"outage": {
  "reason": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
},
"overall_security": "Level 2",
"paths": {
  "initiator": "3a",
  "node.name": "vsim4",
  "node.uuid": "cf7fe057-526d-11ec-af4e-0050568e9df0",
  "port_name": "A",
  "port_type": "sas",
  "vmdisk_hypervisor_file_name": "xvds vol10a0567ae156ca59f6",
  "wwnn": "5000c2971c1b2b8c",
  "wwpn": "5000c2971c1b2b8d"
},
"physical_size": 228930,
"pool": "pool10",
"protection_mode": "data",
"rated_life_used_percent": 10,
"right_size_sector_count": 1172123568,
"rpm": 15000,
"sector_count": 1172123568,
"serial_number": "KHG2VX8R",
"shelf": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},

```

```

    "uid": "7777841915827391056"
  },
  "state": "present",
  "stats": {
    "average_latency": 3,
    "iops_total": 12854,
    "path_error_count": 0,
    "power_on_hours": 21016,
    "throughput": 1957888
  },
  "storage_pool": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "storage_pool_1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "type": "ssd",
  "uid":
"002538E5:71B00B2F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000",
  "usable_size": 959934889984,
  "vendor": "NETAPP",
  "virtual": {
    "container": "nviet12122018113936-rg",
    "object": "f1fu63se",
    "storage_account": "nviet12122018113936ps",
    "target_address": "string"
  }
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregates

Name	Type	Description
_links	_links	
name	string	
uuid	string	

dr_node

Name	Type	Description
name	string	
uuid	string	

drawer

Name	Type	Description
id	integer	
slot	integer	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

The message and code detailing the error state of this disk.

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.

disk_error_info

Name	Type	Description
reason	error	The message and code detailing the error state of this disk.
type	string	Disk error type.

home_node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

key_id

Name	Type	Description
data	string	Key ID of the data authentication key
fips	string	Key ID of the FIPS authentication key

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

error

This error message and code explaining the disk failure.

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.

outage

Indicates if a disk has an entry in the failed disk registry, along with the reason for the failure.

Name	Type	Description
persistently_failed	boolean	Indicates whether RAID maintains the state of this disk as failed across reboots.
reason	error	This error message and code explaining the disk failure.

disk_path_info

Name	Type	Description
initiator	string	Initiator port.
node.name	string	Controller with the initiator port for this path.
node.uuid	string	Controller UUID, to identify node for this path.
port_name	string	Name of the disk port.
port_type	string	Disk port type.
vmdisk_hypervisor_file_name	string	Virtual disk hypervisor file name.
wwnn	string	Target device's World Wide Node Name.

Name	Type	Description
wwpn	string	Target device's World Wide Port Name.

shelf

Shelf

Name	Type	Description
_links	_links	
uid	string	

stats

Name	Type	Description
average_latency	integer	Average I/O latency across all active paths, in milliseconds.
iops_total	integer	Total I/O operations per second read and written to this disk across all active paths.
path_error_count	integer	Disk path error count; failed I/O operations.
power_on_hours	integer	Hours powered on.
throughput	integer	Total disk throughput per second across all active paths, in bytes.

storage_pool

Shared Storage Pool

Name	Type	Description
_links	_links	
name	string	
uuid	string	

virtual

Information about backing storage for disks on cloud platforms.

Name	Type	Description
container	string	Container name of the virtual disk.
object	string	Object name of the virtual disk.
storage_account	string	Storage account name of the virtual disk.
target_address	string	Target address of the virtual disk.

error

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

Manage storage file clones

Storage file clone endpoint overview

Overview

You can use these API's to create file clones, retrieve split status and manage split loads. These endpoints are used for cloning files within a volume, without taking much of extra space. Child and parent clones shares the unchanged blocks of data.

A file clone split operation detach child clone from its parent. Split operations use space. To ensure that file clone create operation is not affected by split, file clone tokens are use to reserve space. API endpoints can be used to update the validity and space reserved by token.

File clone APIs

The following APIs are used to perform the following operations:

– POST `/api/storage/file/clone`

– GET `/api/storage/file/clone/split-status`

– PATCH `/api/storage/file/clone/split-loads/{node.uuid}`

– GET /api/storage/file/clone/split-loads/{node.uuid}

– GET /api/storage/file/clone/split-loads

– GET /api/storage/file/clone/tokens/

– DELETE /api/storage/file/clone/tokens/{node.uuid}/{token.uuid}

– PATCH /api/storage/file/clone/tokens/{node.uuid}/{token.uuid}

Create a clone of the file

POST /storage/file/clone

Introduced In: 9.6

Creates a clone of the file.

Required Properties

- `source_path`
- `destination_path`
- `volume.uuid` and `volume.name` - Instance UUID and name of volume in which to create clone.

Optional Properties

- `range` - Required only in the case of a sub file clone.
- `autodelete` - Marks a cloned file for auto deletion.
- `backup` - Cloned file is used as a backup.

Related Ontap commands

- `volume file clone create`

Creating file clones

The POST operation is used to create file clones with the specified attributes in body. Set the `volume.name` and `volume.uuid` to identify the volume.

Set `source_path` and `destination_path` to identify the file path of original and copied file. In case of full file clone, the new file is created using `destination_path`.
 In case of a sub file clone, set `range` in the format `source-file-block-number:destination-file-block-number:block-count`. The API returns an error for the following overlapping conditions: (a) if source and destination files are same and any of the source ranges overlap with any of the destination ranges. (b) if any of the source ranges overlap amongst themselves. (c) if any of the destination ranges overlap amongst themselves. If not provided, full file cloning is assumed.

If set to `autodelete`, the cloned file is deleted when the volumes are full.


```
# The API:
curl -X POST "https://<mgmt_ip>/api/storage/file/clone" -H "accept:
application/hal+json" -d '{"volume": {"name": "vol1", "uuid": "40e0fdc5-
c28f-11eb-8270-005056bbeb0b"}, "source_path": "f1", "destination_path":
"f2_c1"}'

# The response:
{
"job": {
  "uuid": "0d025fd9-c4dc-11eb-adb5-005056bbeb0b",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/0d025fd9-c4dc-11eb-adb5-005056bbeb0b"
    }
  }
}
}
```

Learn More

- [\[DOC /storage/file/clone\]](#)

Parameters

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

Request Body

Name	Type	Description
autodelete	boolean	Mark clone file for auto deletion.
destination_path	string	Relative path of the clone/destination file in the volume.

Name	Type	Description
is_backup	boolean	Mark clone file for backup.
overwrite_destination	boolean	Destination file gets overwritten.
range	array[string]	List of block ranges for sub-file cloning in the format "source-file-block-number:destination-file-block-number:block-count"
source_path	string	Relative path of the source file in the volume.
token_uuid	string	UUID of existing clone token with reserved split load.
volume	volume	

Example request

```
{
  "destination_path": "dest_file1, dir1/dest_file2",
  "range": [
    36605,
    73210
  ],
  "source_path": "src_file1, dir1/src_file2,
  ../.snapshot/snap1/src_file3",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```

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

```

Headers

Name	Description	Type
Location	Useful for tracking the resource location	string

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

volume

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

file_clone

File clone

Name	Type	Description
autodelete	boolean	Mark clone file for auto deletion.
destination_path	string	Relative path of the clone/destination file in the volume.
is_backup	boolean	Mark clone file for backup.
overwrite_destination	boolean	Destination file gets overwritten.

Name	Type	Description
range	array[string]	List of block ranges for sub-file cloning in the format "source-file-block-number:destination-file-block-number:block-count"
source_path	string	Relative path of the source file in the volume.
token_uuid	string	UUID of existing clone token with reserved split load.
volume	volume	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve the clone split load of a node

GET /storage/file/clone/split-loads

Introduced In: 9.10

Retrieves the clone split load of a node.

Related Ontap Commands

- `volume file clone split load show`

Retrieving file clone split load related information

The GET operation can be used to retrieve information about clone split load data. Split load data is the data currently undergoing the split. There is a limit on split load data. This API communicates how much data is undergoing split and how much can still be processed.


```
# The API:
/api/storage/file/clone/split-loads

# The call:
curl -X GET "https://<mgmt_ip>/api/storage/file/clone/split-loads" -H
"accept: application/hal+json"

# The response:
{
  "records": [
    {
      "node": {
        "uuid": "158d592f-a829-11eb-a47b-005056bb46d7",
        "name": "node1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/158d592f-a829-11eb-a47b-
005056bb46d7"
          }
        }
      },
      "load": {
        "maximum": 35184372088832,
        "current": 0,
        "token_reserved": 0,
        "allowable": 35184372088832
      },
      "_links": {
        "self": {
          "href": "/api/storage/file/clone/split-loads/158d592f-a829-11eb-
```

```

a47b-005056bb46d7"
  }
}
},
{
  "node": {
    "uuid": "9686b8d1-a828-11eb-80d8-005056bbe7b6",
    "name": "node2",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/9686b8d1-a828-11eb-80d8-005056bbe7b6"
      }
    }
  },
  "load": {
    "maximum": 35184372088832,
    "current": 0,
    "token_reserved": 0,
    "allowable": 35184372088832
  },
  "_links": {
    "self": {
      "href": "/api/storage/file/clone/split-loads/9686b8d1-a828-11eb-80d8-005056bbe7b6"
    }
  }
}
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/storage/file/clone/split-loads"
  }
}
}
}

```

Learn More

- [\[DOC /storage/file/clone\]](#)

Parameters

Name	Type	In	Required	Description
node.uuid	string	query	False	Filter by node.uuid

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

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

Response

Status: 200, Ok

Name	Type	Description
_links	collection_links	
num_records	integer	Number of records.
records	array[split_load]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "load": {
      "allowable": 0,
      "current": 0,
      "token_reserved": 0
    },
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

collection_links

Name	Type	Description
next	href	
self	href	

self_link

Name	Type	Description
self	href	

load

Name	Type	Description
allowable	integer	Specifies the available file clone split load on the node.
current	integer	Specifies the current on-going file clone split load on the node.
maximum	integer	Specifies the maximum allowable file clone split load on the node at any point in time.
token_reserved	integer	Specifies the file clone split load on the node reserved for tokens.

_links

Name	Type	Description
self	href	

node_reference

Name	Type	Description
_links	_links	
name	string	

Name	Type	Description
uuid	string	

split_load

Name	Type	Description
_links	self_link	
load	load	
node	node_reference	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve the volume file clone split load

GET /storage/file/clone/split-loads/{node.uuid}

Introduced In: 9.10

Retrieve Volume File Clone Split Load REST

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Node Uuid

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

Response

Status: 200, Ok

Name	Type	Description
_links	self_link	
load	load	
node	node_reference	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "load": {
    "allowable": 0,
    "current": 0,
    "token_reserved": 0
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

Name	Type	Description
self	href	

load

Name	Type	Description
allowable	integer	Specifies the available file clone split load on the node.
current	integer	Specifies the current on-going file clone split load on the node.
maximum	integer	Specifies the maximum allowable file clone split load on the node at any point in time.
token_reserved	integer	Specifies the file clone split load on the node reserved for tokens.

_links

Name	Type	Description
self	href	

node_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

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

Update the maximum split load

PATCH /storage/file/clone/split-loads/{node.uuid}

Introduced In: 9.10

Updates the maximum split load.

Related Ontap command

- `volume file clone split load modify`

Learn More

- [\[DOC /storage/file/clone\]](#)

```
# The call:
curl -X PATCH "https://<mgmt_IP>/api/storage/file/clone/split-loads/9686b8d1-a828-11eb-80d8-005056bbe7b6" -d '{"load": {"maximum": "16TB" } }'
```

```
# The response to successful patch is empty body
```

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Node UUID

Request Body

Name	Type	Description
_links	self_link	
load	load	
node	node_reference	

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "load": {
    "allowable": 0,
    "current": 0,
    "token_reserved": 0
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

Response

Status: 200, Ok

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

Name	Type	Description
self	href	

load

Name	Type	Description
allowable	integer	Specifies the available file clone split load on the node.
current	integer	Specifies the current on-going file clone split load on the node.
maximum	integer	Specifies the maximum allowable file clone split load on the node at any point in time.
token_reserved	integer	Specifies the file clone split load on the node reserved for tokens.

_links

Name	Type	Description
self	href	

node_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

split_load

Name	Type	Description
_links	self_link	

Name	Type	Description
load	load	
node	node_reference	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve file clone split status of all volumes in the node

GET /storage/file/clone/split-status

Introduced In: 9.10

Retrieves file clone split status of all volumes in the node.

Learn More

- [\[DOC /storage/file/clone\]](#)

```
# The API:
/api/storage/file/clone/split-status

# The call:
curl -X GET "https://<mgmt_ip>/api/storage/file/clone/split-status" -H
"accept: application/hal+json"

# The response:
{
```

```
"records": [
  {
    "volume": {
      "uuid": "ac559964-57a3-40cf-b5cb-f3cb99151a7d",
      "name": "voll1",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/ac559964-57a3-40cf-b5cb-
f3cb99151a7d"
        }
      }
    },
    "svm": {
      "name": "vs1"
    },
    "pending_splits": 0,
    "unsplit_clone_size": 0,
    "_links": {
      "self": {
        "href": "/api/storage/file/clone/split-status/ac559964-57a3-40cf-
b5cb-f3cb99151a7d"
      }
    }
  },
  {
    "volume": {
      "uuid": "32d95d48-d8b7-11eb-a41d-005056bb3837",
      "name": "vs1_root",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/32d95d48-d8b7-11eb-a41d-
005056bb3837"
        }
      }
    },
    "svm": {
      "name": "vs1"
    },
    "pending_splits": 0,
    "unsplit_clone_size": 0,
    "_links": {
      "self": {
        "href": "/api/storage/file/clone/split-status/32d95d48-d8b7-11eb-
a41d-005056bb3837"
      }
    }
  }
]
```

```

    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/storage/file/clone/split-status"
    }
  }
}
}

```

Parameters

Name	Type	In	Required	Description
pending_splits	integer	query	False	Filter by pending_splits
unsplit_clone_size	integer	query	False	Filter by unsplit_clone_size
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
volume.uuid	string	query	False	Filter by volume.uuid
volume.name	string	query	False	Filter by volume.name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • 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	collection_links	
num_records	integer	Number of Records
records	array[split_status]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "pending_splits": 0,
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "unsplit_clone_size": 0,
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

collection_links

Name	Type	Description
next	href	
self	href	

self_link

Name	Type	Description
self	href	

_links

Name	Type	Description
self	href	

svm

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

volume

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

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

split_status

Name	Type	Description
_links	self_link	
pending_splits	integer	Specifies the number of pending file clone split operations in the volume.
svm	svm	
unsplit_clone_size	integer	Specifies the space occupied by unsplit file clones in the volume.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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

Retrieve file clone split status of all volumes

GET /storage/file/clone/split-status/{volume.uuid}

Introduced In: 9.10

Retrieves file clone split status of all volumes in the node.

Related Ontap commands

- `volume file clone split status`

Learn More

- [\[DOC /storage/file/clone\]](#)

Retrieves the information of split status.

The GET operation retrieves information about split processes in the volume.
 `pending-clone-splits` is the number of files for which file clone split is not yet completed.
 `unsplit-size` is the sum of all sizes, in bytes in the volume that is not split.


```

# The API:
/api/storage/file/clone/split-status/{volume.uuid}

# The call:
curl -X GET "https://<mgmt_ip>/api/storage/file/clone/split-
status/ac559964-57a3-40cf-b5cb-f3cb99151a7d" -H "accept:
application/hal+json"

# The response:
{
  "volume": {
    "uuid": "ac559964-57a3-40cf-b5cb-f3cb99151a7d",
    "name": "voll1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/ac559964-57a3-40cf-b5cb-f3cb99151a7d"
      }
    }
  },
  "svm": {
    "name": "vs1"
  },
  "pending_splits": 0,
  "unsplit_clone_size": 0,
  "_links": {
    "self": {
      "href": "/api/storage/file/clone/split-status/ac559964-57a3-40cf-b5cb-
f3cb99151a7d"
    }
  }
}

```

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume Instance UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	self_link	
pending_splits	integer	Specifies the number of pending file clone split operations in the volume.
svm	svm	
unsplit_clone_size	integer	Specifies the space occupied by unsplit file clones in the volume.
volume	volume	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "pending_splits": 0,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unsplit_clone_size": 0,
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

Name	Type	Description
self	href	

_links

Name	Type	Description
self	href	

svm

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

volume

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve information for a token

GET `/storage/file/clone/tokens`

Introduced In: 9.10

Retrieves information for the specified token.

Related Ontap command

- `volume file clone token show`

Learn More

- [\[DOC /storage/file/clone\]](#)

Retrieving information on clone tokens

```

# The API:
/api/storage/file/clone/tokens

# The call:
curl -X GET "https://<mgmt_ip>/api/storage/file/clone/tokens" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "node": {
        "uuid": "97255711-a1ad-11eb-92b2-0050568eb2ca",
        "name": "node1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/97255711-a1ad-11eb-92b2-
0050568eb2ca"
          }
        }
      },
      "uuid": "905c42ce-a74b-11eb-bd86-0050568ec7ae",
      "reserve_size": 10240,
      "expiry_time": {
        "limit": "PT1H10M",
        "left": "PT1H9M"
      },
      "_links": {
        "self": {
          "href": "/api/storage/file/clone/tokens/97255711-a1ad-11eb-92b2-
0050568eb2ca/905c42ce-a74b-11eb-bd86-0050568ec7ae"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/storage/file/clone/tokens"
    }
  }
}

```

Parameters

Name	Type	In	Required	Description
uuid	string	query	False	Filter by uuid
reserve_size	integer	query	False	Filter by reserve_size
node.uuid	string	query	False	Filter by node.uuid
node.name	string	query	False	Filter by node.name
expiry_time.left	string	query	False	Filter by expiry_time.left
expiry_time.limit	string	query	False	Filter by expiry_time.limit
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • 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	collection_links	
num_records	integer	Number of records.
records	array[token]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

collection_links

Name	Type	Description
next	href	
self	href	

self_link

Name	Type	Description
self	href	

expiry_time

Name	Type	Description
left	string	Specifies the time remaining before the given token expires in ISO-8601 format.
limit	string	Specifies when the given token expires in ISO-8601 format.

_links

Name	Type	Description
self	href	

node_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

token

token

Name	Type	Description
_links	self_link	
expiry_time	expiry_time	
node	node_reference	
reserve_size	integer	Specifies the available reserve in the file clone split load for the given token.
uuid	string	Token UUID.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Create a new token to reserve the split load

POST /storage/file/clone/tokens

Introduced In: 9.10

Creates a new token to reserve the split load.

Required Properties

- `node.uuid`
- `reserve-size`

Optional Properties

- `expiry_time.limit`
- `expiry_time.left`

Related Ontap Commands

- `volume file clone token create`

Learn More

- [DOC /storage/file/clone]

Creating clone tokens to reserve space for clone creation on the node

There is a limit on the amount of clone data that can undergo a split at a point of time on the node (clone split load). Clone tokens are used to reserve space from clone split load for clone creation. The POST operation is used to create clone tokens with `reserve-size` and `expiry-time.limit` in the body.


```
# The API
/api/storage/file/clone/tokens

# The call
curl -X POST "https://<mgmt_ip>/api/storage/file/clone/tokens" -H "accept:
application/hal+json" -d '{"node": {"uuid": "97255711-a1ad-11eb-92b2-
0050568eb2ca"}, "reserve_size": "40M", "expiry_time": { "limit": "4200"
} }'

# The response
{
  "num_records": 1,
  "records": [
    {
      "node": {
        "name": "node1"
      },
      "uuid": "286f6ae4-c94d-11eb-adb5-005056bbeb0b",
      "reserve_size": 41943040,
      "_links": {
        "self": {
          "href": "/api/storage/file/clone/tokens/97255711-a1ad-11eb-92b2-
0050568eb2ca"
        }
      }
    }
  ]
}
```

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
_links	self_link	
expiry_time	expiry_time	
node	node_reference	
reserve_size	integer	Specifies the available reserve in the file clone split load for the given token.
uuid	string	Token UUID.

Response

Status: 201, Created

Name	Type	Description
_links	collection_links	
num_records	integer	Number of records.
records	array[token]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourceLink"
    },
    "self": {
      "href": "/api/resourceLink"
    }
  },
  "num_records": 1,
  "records": {
  }
}
```

Headers

Name	Description	Type
Location	Useful for tracking the resource location	string

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

Name	Type	Description
self	href	

expiry_time

Name	Type	Description
left	string	Specifies the time remaining before the given token expires in ISO-8601 format.
limit	string	Specifies when the given token expires in ISO-8601 format.

_links

Name	Type	Description
self	href	

node_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

token

token

Name	Type	Description
_links	self_link	
expiry_time	expiry_time	
node	node_reference	

Name	Type	Description
reserve_size	integer	Specifies the available reserve in the file clone split load for the given token.
uuid	string	Token UUID.

collection_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Delete a specific file clone token

```
DELETE /storage/file/clone/tokens/{node.uuid}/{uuid}
```

Introduced In: 9.10

Deletes a specific file clone token.

Related Ontap command

- `volume file clone token delete`

Delete specific clone token.

```
# The API:
/api/storage/file/clone/tokens/{node.uuid}/{token.uuid}

# The call:
curl -X DELETE "https://<mgmt_ip>/api/storage/file/clone/tokens/97255711-
a1ad-11eb-92b2-0050568eb2ca/909c42ce-a74b-11eb-bd86-0050568ec7ae"

# The successful response is empty body.
```

Learn More

- [DOC /storage/file/clone]

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Node UUID
uuid	string	path	True	Token UUID

Response

Status: 200, Ok

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a file clone token

GET /storage/file/clone/tokens/{node.uuid}/{uuid}

Introduced In: 9.10

Retrieves a file clone token

Related Ontap command

- `volume file clone token show`

Retrieve information for single token.

```
# The call:
curl -X GET "https://<mgmt_ip>/api/storage/file/clone/tokens/97255711-
a1ad-11eb-92b2-0050568eb2ca/905c42ce-a74b-11eb-bd86-0050568ec7ae"

# The response:
{
  "node": {
    "uuid": "97255711-a1ad-11eb-92b2-0050568eb2ca",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/97255711-a1ad-11eb-92b2-0050568eb2ca"
      }
    }
  },
  "uuid": "905c42ce-a74b-11eb-bd86-0050568ec7ae",
  "reserve_size": 41943040,
  "expiry_time": {
    "limit": "PT1H10M",
    "left": "PT1H9M"
  },
  "_links": {
    "self": {
      "href": "/api/storage/file/clone/tokens/97255711-a1ad-11eb-92b2-
0050568eb2ca/905c42ce-a74b-11eb-bd86-0050568ec7ae"
    }
  }
}
```

Learn More

- [\[DOC /storage/file/clone\]](#)

Parameters

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

Response

Status: 200, Ok

Name	Type	Description
_links	self_link	
expiry_time	expiry_time	
node	node_reference	
reserve_size	integer	Specifies the available reserve in the file clone split load for the given token.
uuid	string	Token UUID.

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

Name	Type	Description
self	href	

expiry_time

Name	Type	Description
left	string	Specifies the time remaining before the given token expires in ISO-8601 format.
limit	string	Specifies when the given token expires in ISO-8601 format.

_links

Name	Type	Description
self	href	

node_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Update a file clone token

PATCH /storage/file/clone/tokens/{node.uuid}/{uuid}

Introduced In: 9.10

Updates a file clone token.

Related Ontap commands

- `volume file clone token modify`

Modify clone token

Use the PATCH API to update the expiry time associated with the clone token.


```
# The call:
curl -X PATCH "https://<mgmt_ip>/api/storage/file/clone/tokens/97255711-
alad-11eb-92b2-0050568eb2ca/905c42ce-a74b-11eb-bd86-0050568ec7ae" -d
 '{"expiry_time": {"limit": "5400"} }'

# The response for successful PATCH is empty.
```

Learn More

- [\[DOC /storage/file/clone\]](#)

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Node UUID
uuid	string	path	True	Token UUID

Request Body

Name	Type	Description
_links	self_link	
expiry_time	expiry_time	
node	node_reference	
reserve_size	integer	Specifies the available reserve in the file clone split load for the given token.
uuid	string	Token UUID.

Response

Status: 200, Ok

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

Name	Type	Description
self	href	

expiry_time

Name	Type	Description
left	string	Specifies the time remaining before the given token expires in ISO-8601 format.
limit	string	Specifies when the given token expires in ISO-8601 format.

_links

Name	Type	Description
self	href	

node_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

token

token

Name	Type	Description
_links	self_link	
expiry_time	expiry_time	
node	node_reference	

Name	Type	Description
reserve_size	integer	Specifies the available reserve in the file clone split load for the given token.
uuid	string	Token UUID.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Start a file copy operation

POST /storage/file/copy

Introduced In: 9.8

Starts a file copy operation. Only supported on flexible volumes.

Required properties

- `files_to_copy` - List of files with the destination they are to be copied to.

Default property values

- `cutover_time` - `10`
- `hold_quiescence` - `false`
- `max_throughput` - `0`

- `reference_cutover_time - 10`

Related ONTAP commands

- `volume file copy start`

Examples

Copying two files

The POST request is used to copy file(s).

```
# The API:
/api/storage/file/copy

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/file/copy" -H "accept:
application/hal+json" -d
'{"files_to_copy":[{"source":{"volume":{"name":"vol_a"},"svm":{"name":"vs0
"},"path":"d1/src_f1"},"destination":{"volume":{"name":"vol_a"},"svm":{"na
me":"vs0"},"path":"d1/dst_f1"}},
{"source":{"volume":{"name":"vol_a"},"svm":{"name":"vs0"},"path":"d1/src_f
2"},"destination":{"volume":{"name":"vol_a"},"svm":{"name":"vs0"},"path":"
d1/dst_f2"}}}]}'

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

Parameters

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

Request Body

Name	Type	Description
cutover_time	integer	The maximum amount of time (in seconds) that the source can be quiesced before a destination file must be made available for read-write traffic.

Name	Type	Description
files_to_copy	array[files_to_copy]	A list of source files along with the destinations they are copied to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory.
hold_quiescence	boolean	Specifies whether the source file should be held quiescent for the duration of the copy operation.
max_throughput	integer	The maximum amount of data (in bytes) that can be transferred per second in support of this operation.
reference_cutover_time	integer	The maximum amount of time (in seconds) that the source reference file can be quiesced before the corresponding destination file must be made available for read-write traffic.
reference_file	reference_file	

Example request

```
{
  "cutover_time": 10,
  "files_to_copy": {
    "destination": {
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volumel",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    },
    "source": {
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volumel",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  },
}
```

```

"reference_cutover_time": 10,
"reference_file": {
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```

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

```

Headers

Name	Description	Type
Location	Useful for tracking the resource location	string

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
7012352	File locations are inconsistent. All files must be on the same volume.
7012353	Exceeded the file operations supported number of files.
7012354	Unable to pair the number of source files to destination files.
7012357	Cannot start a file operation until all cluster nodes support the file operations capability.
7012358	The specified source path is invalid.
7012359	The specified destination path is invalid.
7012360	The SVMs are not in an intracluster peering relationship.
7012361	The SVMs peering relationship does not include application "file-copy".
7012362	The SVMs are not yet in a peered state yet.
7012363	Cannot copy files. All file operations must be managed by the destination SVM's administrator.
7012365	Copying a file between clusters is not supported.
7012367	A reference path may only be specified if multiple source paths are specified.
7012368	The reference path must have a matching source path.
7012371	The reference cutover time exceeds the maximum allowable time.
7012374	Source volume and destination volume have different home clusters.
7012376	Operation not allowed on a volume that is part of a SnapMirror Synchronous relationship.
7012377	Cannot start a file copy operation on the volume because an active volume conversion is in progress.
7018877	Maximum combined total (50) of file and LUN copy and move operations reached. When one or more of the operations has completed, try the command again.

Error Code	Description
13107223	Operation not supported for FlexGroup volumes or FlexGroup constituents.
196608143	Cannot start operation. The volume is undergoing a secure purge operation.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

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

volume

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

file_reference

Name	Type	Description
path	string	Path of the file or directory.
svm	svm	

Name	Type	Description
volume	volume	

files_to_copy

Name	Type	Description
destination	file_reference	
source	file_reference	

reference_file

Name	Type	Description
path	string	The source reference file. If a reference file is specified, data for other files being copied will be transferred as a difference from the reference file. This can save bandwidth and destination storage if the specified source files share blocks. If provided, this input must match one of the source file paths. This input need not be provided if only one source file is specified.
volume	volume	

file_copy

File copy

Name	Type	Description
cutover_time	integer	The maximum amount of time (in seconds) that the source can be quiesced before a destination file must be made available for read-write traffic.
files_to_copy	array[files_to_copy]	A list of source files along with the destinations they are copied to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory.
hold_quiescence	boolean	Specifies whether the source file should be held quiescent for the duration of the copy operation.

Name	Type	Description
max_throughput	integer	The maximum amount of data (in bytes) that can be transferred per second in support of this operation.
reference_cutover_time	integer	The maximum amount of time (in seconds) that the source reference file can be quiesced before the corresponding destination file must be made available for read-write traffic.
reference_file	reference_file	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage storage file moves

Storage file moves endpoint overview

Overview

You can use this API to start a file move operation between two FlexVol volumes or within a FlexGroup volume, and view the status of all on-going file move operations in the cluster.

– GET /api/storage/file/moves

– GET /api/storage/file/moves/{node.uuid}/{uuid}/{index}

– POST /api/storage/file/moves

Examples

Moving two files from one FlexVol volume to the other FlexVol volume

```
# The API:
/api/storage/file/moves

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/file/moves" -H "accept:
application/hal+json" -d @move_between_fv.txt
move_between_fv.txt:
{
  "files_to_move":{
    "sources":[
      {
        "svm":{
          "name":"vs0"
        },
        "volume":{
          "name":"fv1"
        },
        "path":"dir1/f1.txt"
      },
      {
        "svm":{
          "name":"vs0"
        },
        "volume":{
          "name":"fv1"
        },
        "path":"dir1/f2.txt"
      }
    ]
  }
}
```

```
"destinations":[
  {
    "svm":{
      "name":"vs0"
    },
    "volume":{
      "name":"fv2"
    },
    "path":"dir2/f1.txt"
  },
  {
    "svm":{
      "name":"vs0"
    },
    "volume":{
      "name":"fv2"
    },
    "path":"dir2/f2.txt"
  }
]
}
}

# The response:
{ }
```

Moving two files from one FlexVol volume to the other FlexVol volume (only specifying the destination directory)

```

# The API:
/api/storage/file/moves

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/file/moves" -H "accept:
application/hal+json" -d @move_between_fv_dir.txt
move_between_fv_dir.txt:
{
"files_to_move":{
  "sources":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv1"
      },
      "path":"dir1/f3.txt"
    },
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv1"
      },
      "path":"dir1/f4.txt"
    }
  ],
  "destinations":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv2"
      },
      "path":"dir2/"
    }
  ]
}
}

# The response:
{ }

```

Moving multiple files from one FlexVol volume to the other FlexVol volume and providing a source reference file

```
# The API:
/api/storage/file/moves

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/file/moves" -H "accept:
application/hal+json" -d @move_between_fv_source.txt
move_between_fv_source.txt:
{
"files_to_move":{
  "sources":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv1"
      },
      "path":"dir1/f5.txt"
    },
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv1"
      },
      "path":"dir1/f6.txt"
    },
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fv1"
      },
      "path":"dir1/f7.txt"
    }
  ],
  "destinations":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
```

```

    "name": "fv2"
  },
  "path": "dir2/f5.txt"
},
{
  "svm": {
    "name": "vs0"
  },
  "volume": {
    "name": "fv2"
  },
  "path": "dir2/f6.txt"
},
{
  "svm": {
    "name": "vs0"
  },
  "volume": {
    "name": "fv2"
  },
  "path": "dir2/f700.txt"
}
]
},
"reference": {
  "svm": {
    "name": "vs0"
  },
  "volume": {
    "name": "fv1"
  },
  "path": "dir1/f6.txt"
}
}

# The response:
{ }

```

Moving a file between two FlexGroup volume constituents in the same FlexGroup volume


```

# The API:
/api/storage/file/moves

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/file/moves" -H "accept:
application/hal+json" -d @move_between_fg.txt
move_between_fg.txt:
{
"files_to_move":{
  "sources":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fg2"
      },
      "path":"test/file.txt"
    }
  ],
  "destinations":[
    {
      "volume":{
        "name":"fg2__0008"
      }
    }
  ]
}
}

# The response:
{ }

```

Automatically selecting a destination constituent to move a file in a FlexGroup volume for capacity rebalancing

```
# The API:
/api/storage/file/moves

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/file/moves?automatic=true" -H
"accept: application/hal+json" -d @move_between_fg_automatic.txt
move_between_fg_automatic.txt:
{
"files_to_move":{
  "sources":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fg2"
      },
      "path":"test/file2.txt"
    }
  ]
}
}

# The response:
{ }
```

Moving a file between two FlexGroup volume constituents in the same FlexGroup volume using the "force" and "disruptive" options

```

# The API:
/api/storage/file/moves

# The call:
curl -X POST "https://<mgmt-
ip>/api/storage/file/moves?force=true&disruptive=true" -H "accept:
application/hal+json" -d @move_between_fg_force.txt
move_between_fg_force.txt
{
"files_to_move":{
  "sources":[
    {
      "svm":{
        "name":"vs0"
      },
      "volume":{
        "name":"fg2"
      },
      "path":"test/file3.txt"
    }
  ],
  "destinations":[
    {
      "volume":{
        "name":"fg2__0008"
      }
    }
  ]
}
}

# The response:
{ }

```

Retrieving file move operations

```

# The API:
/api/storage/file/moves

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

# The response:
{

```

```

"records": [
  {
    "node": {
      "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-005056acf2dd"
        }
      }
    },
    "uuid": "e12bc78d-36bb-4274-8163-fb8c21d59c9b",
    "index": 0,
    "source": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "4e919b6d-1c76-11ec-8e1b-005056acf2dd",
        "name": "fg2__0008"
      },
      "path": "test/file2.txt"
    },
    "destination": {
      "svm": {
        "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
        "name": "vs0"
      },
      "volume": {
        "uuid": "4d14f2f6-1c76-11ec-8e1b-005056acf2dd",
        "name": "fg2__0005"
      },
      "path": "test/file2.txt"
    },
    "_links": {
      "self": {
        "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-005056acf2dd/e12bc78d-36bb-4274-8163-fb8c21d59c9b/0"
      }
    }
  },
  {
    "node": {
      "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",

```

```

    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-
005056acf2dd"
      }
    }
  },
  "uuid": "ce2af347-586d-4b31-b728-1e925f51fdfc",
  "index": 1,
  "source": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    "volume": {
      "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
      "name": "fv1"
    },
    "path": "dir1/f2.txt"
  },
  "destination": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    "volume": {
      "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
      "name": "fv2"
    },
    "path": "dir2/f2.txt"
  },
  "_links": {
    "self": {
      "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/ce2af347-586d-4b31-b728-1e925f51fdfc/1"
    }
  }
},
{
  "node": {
    "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-

```

```

005056acf2dd"
  }
}
},
"uuid": "ce2af347-586d-4b31-b728-1e925f51fdcf",
"index": 0,
"source": {
  "svm": {
    "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
    "name": "vs0"
  },
  "volume": {
    "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
    "name": "fv1"
  },
  "path": "dir1/f1.txt"
},
"destination": {
  "svm": {
    "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
    "name": "vs0"
  },
  "volume": {
    "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
    "name": "fv2"
  },
  "path": "dir2/f1.txt"
},
"_links": {
  "self": {
    "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-
005056acf2dd/ce2af347-586d-4b31-b728-1e925f51fdcf/0"
  }
}
},
{
  "node": {
    "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-
005056acf2dd"
      }
    }
  }
},

```

```

"uuid": "6d12601b-5377-43bf-99f0-b4bec37565e2",
"index": 0,
"source": {
  "svm": {
    "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
    "name": "vs0"
  },
  "volume": {
    "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
    "name": "fv1"
  },
  "path": "dir1/f3.txt"
},
"destination": {
  "svm": {
    "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
    "name": "vs0"
  },
  "volume": {
    "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
    "name": "fv2"
  },
  "path": "dir2/f3.txt"
},
"_links": {
  "self": {
    "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-005056acf2dd/6d12601b-5377-43bf-99f0-b4bec37565e2/0"
  }
}
},
{
  "node": {
    "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-005056acf2dd"
      }
    }
  },
  "uuid": "6d12601b-5377-43bf-99f0-b4bec37565e2",
  "index": 1,
  "source": {
    "svm": {

```

```

    "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
    "name": "vs0"
  },
  "volume": {
    "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
    "name": "fv1"
  },
  "path": "dir1/f4.txt"
},
"destination": {
  "svm": {
    "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
    "name": "vs0"
  },
  "volume": {
    "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
    "name": "fv2"
  },
  "path": "dir2/f4.txt"
},
"_links": {
  "self": {
    "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-005056acf2dd/6d12601b-5377-43bf-99f0-b4bec37565e2/1"
  }
}
},
{
  "node": {
    "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-005056acf2dd"
      }
    }
  },
  "uuid": "bbfdface-0d46-4f5f-9624-72f4869eba81",
  "index": 0,
  "source": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    "volume": {

```



```

    "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
    "name": "fv1"
  },
  "path": "dir1/f5.txt"
},
"destination": {
  "svm": {
    "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
    "name": "vs0"
  },
  "volume": {
    "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
    "name": "fv2"
  },
  "path": "dir2/f5.txt"
},
"_links": {
  "self": {
    "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-005056acf2dd/bbfdface-0d46-4f5f-9624-72f4869eba81/0"
  }
}
},
{
  "node": {
    "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-005056acf2dd"
      }
    }
  },
  "uuid": "bbfdface-0d46-4f5f-9624-72f4869eba81",
  "index": 2,
  "source": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    "volume": {
      "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
      "name": "fv1"
    },
    "path": "dir1/f7.txt"
  }
}

```

```

},
"destination": {
  "svm": {
    "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
    "name": "vs0"
  },
  "volume": {
    "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
    "name": "fv2"
  },
  "path": "dir2/f700.txt"
},
"_links": {
  "self": {
    "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-005056acf2dd/bbfdface-0d46-4f5f-9624-72f4869eba81/2"
  }
}
},
{
  "node": {
    "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-005056acf2dd"
      }
    }
  },
  "uuid": "bbfdface-0d46-4f5f-9624-72f4869eba81",
  "index": 1,
  "source": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    "volume": {
      "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
      "name": "fv1"
    },
    "path": "dir1/f6.txt"
  },
  "destination": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",

```

```

    "name": "vs0"
  },
  "volume": {
    "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
    "name": "fv2"
  },
  "path": "dir2/f6.txt"
},
"_links": {
  "self": {
    "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-005056acf2dd/bbfdface-0d46-4f5f-9624-72f4869eba81/1"
  }
}
},
{
  "node": {
    "uuid": "780255d2-10aa-11ec-a308-005056acf86d",
    "name": "node2",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/780255d2-10aa-11ec-a308-005056acf86d"
      }
    }
  },
  "uuid": "6591a42a-4ea2-4d40-bfb4-38959f6bd68e",
  "index": 0,
  "source": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    "volume": {
      "uuid": "438731dd-1c76-11ec-8e1b-005056acf2dd",
      "name": "fg2__0001"
    },
    "path": "test/file.txt"
  },
  "destination": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    "volume": {
      "uuid": "4e919b6d-1c76-11ec-8e1b-005056acf2dd",

```

```

    "name": "fg2__0008"
  },
  "path": "/test/file.txt"
},
"_links": {
  "self": {
    "href": "/api/storage/file/moves/780255d2-10aa-11ec-a308-005056acf86d/6591a42a-4ea2-4d40-bfb4-38959f6bd68e/0"
  }
}
},
{
  "node": {
    "uuid": "780255d2-10aa-11ec-a308-005056acf86d",
    "name": "node2",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/780255d2-10aa-11ec-a308-005056acf86d"
      }
    }
  },
  "uuid": "1a94e95a-346e-4eb3-969a-110e275cbf18",
  "index": 0,
  "source": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    "volume": {
      "uuid": "438731dd-1c76-11ec-8e1b-005056acf2dd",
      "name": "fg2__0001"
    },
    "path": "test/file3.txt"
  },
  "destination": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    "volume": {
      "uuid": "4e919b6d-1c76-11ec-8e1b-005056acf2dd",
      "name": "fg2__0008"
    },
    "path": "/test/file3.txt"
  },
}

```

```

    "_links": {
      "self": {
        "href": "/api/storage/file/moves/780255d2-10aa-11ec-a308-005056acf86d/1a94e95a-346e-4eb3-969a-110e275cbf18/0"
      }
    }
  },
  "num_records": 10,
  "_links": {
    "self": {
      "href": "/api/storage/file/moves"
    }
  }
}

```

Retrieving all moves in a file move operation

```

# The API:
/api/storage/file/moves/{node.uuid}/{uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/file/moves/*/ce2af347-586d-4b31-b728-1e925f51fdcf/*" -H 'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "node": {
        "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
        "name": "node1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-005056acf2dd"
          }
        }
      },
      "uuid": "ce2af347-586d-4b31-b728-1e925f51fdcf",
      "index": 1,
      "source": {
        "svm": {
          "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
          "name": "vs0"
        }
      }
    }
  ]
}

```

```

    },
    "volume": {
      "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
      "name": "fv1"
    },
    },
    "path": "dir1/f2.txt"
  },
  "destination": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    },
    "volume": {
      "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
      "name": "fv2"
    },
    },
    "path": "dir2/f2.txt"
  },
  "_links": {
    "self": {
      "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-005056acf2dd/ce2af347-586d-4b31-b728-1e925f51fdcf/1"
    }
  }
},
{
  "node": {
    "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-005056acf2dd"
      }
    }
  },
  "uuid": "ce2af347-586d-4b31-b728-1e925f51fdcf",
  "index": 0,
  "source": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    },
    "volume": {
      "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
      "name": "fv1"
    }
  }
}

```

```

    },
    "path": "dir1/f1.txt"
  },
  "destination": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    "volume": {
      "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
      "name": "fv2"
    },
    "path": "dir2/f1.txt"
  },
  "_links": {
    "self": {
      "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-005056acf2dd/ce2af347-586d-4b31-b728-1e925f51fdcf/0"
    }
  }
}
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/storage/file/moves/*/ce2af347-586d-4b31-b728-1e925f51fdcf"
  }
}
}
}

```

Retrieving a specific file move in a file move operation

```

# The API:
/api/storage/file/moves/{node.uuid}/{uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/file/moves/*/ce2af347-586d-4b31-b728-1e925f51fdcf/1" -H 'accept: application/hal+json'

# The response:
{
  "node": {
    "uuid": "76bc12d1-10aa-11ec-a5b5-005056acf2dd",
    "name": "node1",

```

```

    "_links": {
      "self": {
        "href": "/api/cluster/nodes/76bc12d1-10aa-11ec-a5b5-005056acf2dd"
      }
    }
  },
  "uuid": "ce2af347-586d-4b31-b728-1e925f51fdfc",
  "index": 1,
  "source": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    "volume": {
      "uuid": "18fd9110-26f1-11ec-bf0d-005056acf2dd",
      "name": "fv1"
    },
    "path": "dir1/f2.txt"
  },
  "destination": {
    "svm": {
      "uuid": "5b2c8638-10bc-11ec-8e1b-005056acf2dd",
      "name": "vs0"
    },
    "volume": {
      "uuid": "220bdb3a-26f1-11ec-bf0d-005056acf2dd",
      "name": "fv2"
    },
    "path": "dir2/f2.txt"
  },
  "_links": {
    "self": {
      "href": "/api/storage/file/moves/76bc12d1-10aa-11ec-a5b5-005056acf2dd/ce2af347-586d-4b31-b728-1e925f51fdfc/1"
    }
  }
}

```

Retrieve all ongoing file move operations in the cluster

GET /storage/file/moves

Introduced In: 9.11

Retrieves all ongoing file move operations in the cluster.

Related ONTAP commands

- `volume file move show`
- `volume rebalance file-move show`

Parameters

Name	Type	In	Required	Description
source.svm.uuid	string	query	False	Filter by source.svm.uuid
source.svm.name	string	query	False	Filter by source.svm.name
source.path	string	query	False	Filter by source.path
source.volume.uuid	string	query	False	Filter by source.volume.uuid
source.volume.name	string	query	False	Filter by source.volume.name
is_destination_ready	boolean	query	False	Filter by is_destination_ready
max_throughput	integer	query	False	Filter by max_throughput
max_cutover_time	integer	query	False	Filter by max_cutover_time
is_flexgroup	boolean	query	False	Filter by is_flexgroup
scanner.progress	integer	query	False	Filter by scanner.progress
scanner.total	integer	query	False	Filter by scanner.total
scanner.state	string	query	False	Filter by scanner.state
scanner.percent	integer	query	False	Filter by scanner.percent
uuid	string	query	False	Filter by uuid

Name	Type	In	Required	Description
node.uuid	string	query	False	Filter by node.uuid
node.name	string	query	False	Filter by node.name
elapsed_time	integer	query	False	Filter by elapsed_time
failure.arguments.message	string	query	False	Filter by failure.arguments.message
failure.arguments.code	string	query	False	Filter by failure.arguments.code
failure.message	string	query	False	Filter by failure.message
failure.code	string	query	False	Filter by failure.code
failure.target	string	query	False	Filter by failure.target
index	integer	query	False	Filter by index
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
destination.svm.uuid	string	query	False	Filter by destination.svm.uuid
destination.svm.name	string	query	False	Filter by destination.svm.name
destination.path	string	query	False	Filter by destination.path
destination.volume.uuid	string	query	False	Filter by destination.volume.uuid

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

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

Response

Status: 200, Ok

Name	Type	Description
_links	collection_links	
num_records	integer	Number of Records
records	array[file_move]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "cutover_time": 8,
    "destination": {
      "path": "d1/d2/file1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    },
    "elapsed_time": 100,
    "failure": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    },
    "files_to_move": {
```

```

"destinations": {
  "path": "d1/d2/file1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
},
"sources": {
  "path": "d1/d2/file1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
},
"index": 0,
"max_cutover_time": 10,
"max_throughput": 250000,

```

```

"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"reference": {
  "max_cutover_time": 5,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
},
"scanner": {
  "percent": 80,
  "progress": 80000,
  "state": "allocation_map",
  "total": 100000
},
"source": {
  "path": "d1/d2/file1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}

```

```

    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "4fcb3159-a4ee-42b5-bb16-f752f2c430fc",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

collection_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm

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

volume

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

destination

Destination file information.

Name	Type	Description
path	string	
svm	svm	
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Contains the most recent failure reason for move operation.

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.

destinations

Name	Type	Description
path	string	
svm	svm	
volume	volume	

sources

Name	Type	Description
path	string	
svm	svm	
volume	volume	

files_to_move

A list of source files along with the destination file they are moved to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory. This is only used for FlexVol volume file move operations.

Name	Type	Description
destinations	array[destinations]	The destination file information.
sources	array[sources]	The source file information.

node

Node hosting the destination of this move operation.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

SVM of the source reference file.

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

volume

Volume of the source reference file.

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

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

reference

Reference file information.

Name	Type	Description
max_cutover_time	integer	The maximum amount of time, in seconds that the source reference file can be quiesced before the corresponding destination file must be made available for read-write traffic. Not supported in FlexGroup volume file move operations.
path	string	The source reference file. If a reference file is specified, data for other files being moved will be transferred as a difference from the reference file. This can save bandwidth and destination storage if the specified source files share blocks. If provided, this input must match one of the source file paths. This input need not be provided if only one source file is specified. Not supported in FlexGroup volume file move operations.
svm	svm	SVM of the source reference file.
volume	volume	Volume of the source reference file.

scanner

Name	Type	Description
percent	integer	Scanner progress, as a percentage.
progress	integer	Scanner progress, in bytes scanned.
state	string	Status of the file move scanner.
total	integer	Total bytes to be scanned.

source

Source file information.

Name	Type	Description
path	string	
svm	svm	
volume	volume	

svm

SVM of the FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.

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

volume

FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.

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

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

file_move

Details of a file move operation.

Name	Type	Description
cutover_time	integer	Time that the file move operation takes before cutover completes, in seconds.
destination	destination	Destination file information.
elapsed_time	integer	Time elapsed since the start of the file move operation, in seconds.
failure	error	Contains the most recent failure reason for move operation.
files_to_move	files_to_move	A list of source files along with the destination file they are moved to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory. This is only used for FlexVol volume file move operations.

Name	Type	Description
index	integer	An additional unique element identifying one file among many that could possibly be moved as part of a job. File index is an identifier ordered by the file path arrays provided during create. For file move operations that involve only one file, the file-index value of zero is always correct.
is_destination_ready	boolean	Indicates whether the destination file is ready for use.
is_flexgroup	boolean	Indicates whether this is a FlexGroup file move operation.
is_snapshot_fenced	boolean	Indicates whether Snapshot copies are fenced.
max_cutover_time	integer	The maximum amount of time, in seconds that the source can be quiesced before a destination file must be made available for read-write traffic. Not supported in FlexGroup volume file move operations.
max_throughput	integer	Maximum amount of data, in bytes that can be transferred per second in support of this operation. A non-zero value less than 1 MB/s is set to 1 MB/s. A non-zero value greater than 1 MB/s is truncated to the nearest integral megabyte value. If unspecified, the default value is "0" which means no range is set for the data transfer.
node	node	Node hosting the destination of this move operation.
reference	reference	Reference file information.
scanner	scanner	
source	source	Source file information.

Name	Type	Description
svm	svm	SVM of the FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.
uuid	string	The UUID which uniquely identifies the job that started this move operation.
volume	volume	FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.

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.

Start a file move operation

POST /storage/file/moves

Introduced In: 9.11

Starts a file move operation. This API can be used to move files from one FlexVol volume to another FlexVol volume or within a FlexGroup volume for capacity rebalancing. For a FlexGroup volume file move operation, only one source file can be specified in `files_to_move`. The source volume is the FlexGroup volume. The destination volume is the destination FlexGroup volume constituent to move the file to. When `automatic` is true, destination volume is not required. The source path is the path to the file to be moved within the FlexGroup volume. If the destination path is specified, it must be the same as the source path.

Required properties for file move operation

- `files_to_move` - List of files with the destination they are to be moved to.

Optional properties for file move operation

- `reference` - The source reference file for moving multiple files.

Default property values

- `max_throughput` - 0
- `max_cutover_time` - 10
- `reference.max_cutover_time` - 10

Related ONTAP commands

- `volume file move start`
- `volume rebalance file-move start`

Parameters

Name	Type	In	Required	Description
<code>force</code>	boolean	query	False	If true, the FlexGroup volume file move operation breaks the existing lock state on the file being moved. Breaking the lock state may cause a disruption for some client applications. <ul style="list-style-type: none">• Default value:
<code>automatic</code>	boolean	query	False	If true, the FlexGroup volume file move operation selects the destination constituent automatically. <ul style="list-style-type: none">• Default value:

Name	Type	In	Required	Description
disruptive	boolean	query	False	<p>If true, the FlexGroup volume file move operation is disruptive to the clients. The file handle of the file being moved changes. If false, the file handle remains the same. The non-disruptive file move operation is only available on FlexGroup volumes with granular data property enabled.</p> <ul style="list-style-type: none"> • Default value: • Introduced in: 9.12
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
cutover_time	integer	Time that the file move operation takes before cutover completes, in seconds.
destination	destination	Destination file information.
elapsed_time	integer	Time elapsed since the start of the file move operation, in seconds.
failure	error	Contains the most recent failure reason for move operation.

Name	Type	Description
files_to_move	files_to_move	A list of source files along with the destination file they are moved to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory. This is only used for FlexVol volume file move operations.
index	integer	An additional unique element identifying one file among many that could possibly be moved as part of a job. File index is an identifier ordered by the file path arrays provided during create. For file move operations that involve only one file, the file-index value of zero is always correct.
is_destination_ready	boolean	Indicates whether the destination file is ready for use.
is_flexgroup	boolean	Indicates whether this is a FlexGroup file move operation.
is_snapshot_fenced	boolean	Indicates whether Snapshot copies are fenced.
max_cutover_time	integer	The maximum amount of time, in seconds that the source can be quiesced before a destination file must be made available for read-write traffic. Not supported in FlexGroup volume file move operations.
max_throughput	integer	Maximum amount of data, in bytes that can be transferred per second in support of this operation. A non-zero value less than 1 MB/s is set to 1 MB/s. A non-zero value greater than 1 MB/s is truncated to the nearest integral megabyte value. If unspecified, the default value is "0" which means no range is set for the data transfer.
node	node	Node hosting the destination of this move operation.

Name	Type	Description
reference	reference	Reference file information.
scanner	scanner	
source	source	Source file information.
svm	svm	SVM of the FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.
uuid	string	The UUID which uniquely identifies the job that started this move operation.
volume	volume	FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.

Example request

```
{
  "cutover_time": 8,
  "destination": {
    "path": "d1/d2/file1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "elapsed_time": 100,
  "failure": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "files_to_move": {
    "destinations": {
      "path": "d1/d2/file1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  }
}
```

```

    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "sources": {
    "path": "d1/d2/file1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
},
"index": 0,
"max_cutover_time": 10,
"max_throughput": 250000,
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"reference": {

```

```

    "max_cutover_time": 5,
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "scanner": {
    "percent": 80,
    "progress": 80000,
    "state": "allocation_map",
    "total": 100000
  },
  "source": {
    "path": "d1/d2/file1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
}

```



```

},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "4fcb3159-a4ee-42b5-bb16-f752f2c430fc",
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}

```

Response

Status: 201, Created

Headers

Name	Description	Type
Location	Useful for tracking the resource location	string

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
917505	SVM not found.
918236	The specified "volume.uuid" and "volume.name" do not refer to the same volume.

Error Code	Description
2621706	The specified "svm.uuid" and "svm.name" do not refer to the same SVM.
7012352	File locations are inconsistent. All files must be on the same volume.
7012353	Exceeded the file operations supported number of files.
7012354	Unable to pair the number of source files to destination files.
7012357	Cannot start a file operation until all cluster nodes support the file operations capability.
7012358	The specified source path is invalid.
7012359	The specified destination path is invalid.
7012360	The SVMs are not in an intracluster peering relationship.
7012361	The SVMs peering relationship does not include application "file-move".
7012362	The SVMs are not yet in a peered state.
7012363	Cannot move files. All file operations must be managed by the destination SVM's administrator.
7012365	Copying a file between clusters is not supported.
7012367	A reference path may only be specified if multiple source paths are specified.
7012368	The reference path must have a matching source path.
7012371	The reference cutover time exceeds the maximum allowable time.
7012374	Source volume and destination volume have different home clusters.
7012376	Operation not allowed on a volume that is part of a SnapMirror Synchronous relationship.
7012377	Cannot start a file move operation on the volume because an active volume conversion is in progress.
7013352	One or more source files must be specified in "files_to_move.sources".
7013353	One or more destination files or directories must be specified in "files_to_move.destinations".
7013354	"path" must be specified.
7013355	Moving files between FlexVol volumes and FlexGroup volumes or constituents is not supported.

Error Code	Description
7013356	Moving files between FlexVol volumes and FlexGroup constituents is not supported.
7013357	The specified volume could not be found.
7013358	The specified SVM or volume UUID could not be found.
7013359	The SVM and volume must both be provided.
7018877	Maximum combined total (50) of file and LUN copy and move operations reached. When one or more of the operations has completed, try the command again.
7018937	The file is already on the destination constituent.
13107222	Internal error.
13107415	Failed to lookup a volume property.
13107431	Failed to lookup an SVM property.
13109260	Failed to enable granular data on the volume.
144179201	The file move start operation failed.
144179206	Source file does not exist.
144179207	Volume capacity balancing requires an effective cluster version of 9.10.1 or later.
144180200	Destination constituent not a member of FlexGroup volume.
144180201	Destination constituent not properly configured.
144180203	Volume capacity rebalancing is not supported on FlexCache volumes.
144180204	Volume capacity rebalancing is not supported on object store volumes.
144180205	The system is busy.
144180206	File movement with automatic destination constituent selection only supported on FlexGroup volumes with more than one constituent.
144180207	Volume capacity rebalancing is not supported on inactive MetroCluster configurations.
144180208	Disruptive file movement is not supported when granular data is enabled on the volume. Try the operation again using "disruptive=false".
144181200	Too many source or destination files are specified for a file move within a FlexGroup volume. There must be one source file identifying a file on a FlexGroup volume and either zero or one destination files identifying the destination constituent.

Error Code	Description
144181202	For a file move within a FlexGroup volume, the source volume must be a FlexGroup volume, and the destination volume must be a constituent.
144181203	A destination constituent must be provided in "files_to_move.destinations" if it is not being selected automatically. Use the "automatic" query to enable automatic destination constituent selection.
144181204	A destination constituent is provided while automatic destination constituent selection is enabled with the "automatic" query.
144181205	The destination volume is not a constituent. For a file move within a FlexGroup volume, the destination volume must be a constituent of the source FlexGroup volume.
144181207	The destination constituent SVM is not the same as the source SVM. For a file move within a FlexGroup volume, the destination constituent must be a constituent of the source FlexGroup volume.
144181208	The destination file path is different from the source file path. For a file move within a FlexGroup volume, the path of the source file does not change.
144181209	The specified SVM for the destination constituent differs from the SVM of the source FlexGroup volume. For a file move within a FlexGroup volume, the destination constituent must be a constituent of the source FlexGroup volume.
144182201	Volume capacity rebalancing using non-disruptive file move operations and granular data requires an effective cluster version of 9.11.1 or later.
144182206	The "force" parameter is not supported unless the "disruptive" parameter is specified as "true".
196608143	Cannot start the operation. The volume is undergoing a secure purge operation.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

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

volume

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

destination

Destination file information.

Name	Type	Description
path	string	
svm	svm	

Name	Type	Description
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Contains the most recent failure reason for move operation.

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.

destinations

Name	Type	Description
path	string	
svm	svm	
volume	volume	

sources

Name	Type	Description
path	string	
svm	svm	
volume	volume	

files_to_move

A list of source files along with the destination file they are moved to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory. This is only used for FlexVol volume file move operations.

Name	Type	Description
destinations	array[destinations]	The destination file information.
sources	array[sources]	The source file information.

node

Node hosting the destination of this move operation.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

SVM of the source reference file.

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

volume

Volume of the source reference file.

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

reference

Reference file information.

Name	Type	Description
max_cutover_time	integer	The maximum amount of time, in seconds that the source reference file can be quiesced before the corresponding destination file must be made available for read-write traffic. Not supported in FlexGroup volume file move operations.
path	string	The source reference file. If a reference file is specified, data for other files being moved will be transferred as a difference from the reference file. This can save bandwidth and destination storage if the specified source files share blocks. If provided, this input must match one of the source file paths. This input need not be provided if only one source file is specified. Not supported in FlexGroup volume file move operations.
svm	svm	SVM of the source reference file.
volume	volume	Volume of the source reference file.

scanner

Name	Type	Description
percent	integer	Scanner progress, as a percentage.
progress	integer	Scanner progress, in bytes scanned.
state	string	Status of the file move scanner.
total	integer	Total bytes to be scanned.

source

Source file information.

Name	Type	Description
path	string	
svm	svm	
volume	volume	

svm

SVM of the FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.

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

volume

FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.

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

file_move

Details of a file move operation.

Name	Type	Description
cutover_time	integer	Time that the file move operation takes before cutover completes, in seconds.

Name	Type	Description
destination	destination	Destination file information.
elapsed_time	integer	Time elapsed since the start of the file move operation, in seconds.
failure	error	Contains the most recent failure reason for move operation.
files_to_move	files_to_move	A list of source files along with the destination file they are moved to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory. This is only used for FlexVol volume file move operations.
index	integer	An additional unique element identifying one file among many that could possibly be moved as part of a job. File index is an identifier ordered by the file path arrays provided during create. For file move operations that involve only one file, the file-index value of zero is always correct.
is_destination_ready	boolean	Indicates whether the destination file is ready for use.
is_flexgroup	boolean	Indicates whether this is a FlexGroup file move operation.
is_snapshot_fenced	boolean	Indicates whether Snapshot copies are fenced.
max_cutover_time	integer	The maximum amount of time, in seconds that the source can be quiesced before a destination file must be made available for read-write traffic. Not supported in FlexGroup volume file move operations.

Name	Type	Description
max_throughput	integer	Maximum amount of data, in bytes that can be transferred per second in support of this operation. A non-zero value less than 1 MB/s is set to 1 MB/s. A non-zero value greater than 1 MB/s is truncated to the nearest integral megabyte value. If unspecified, the default value is "0" which means no range is set for the data transfer.
node	node	Node hosting the destination of this move operation.
reference	reference	Reference file information.
scanner	scanner	
source	source	Source file information.
svm	svm	SVM of the FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.
uuid	string	The UUID which uniquely identifies the job that started this move operation.
volume	volume	FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.

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.

Check the status of an active file move operation

GET /storage/file/moves/{node.uuid}/{uuid}/{index}

Introduced In: 9.12

Overview

Retrieve the status of an on-going file move operation.

Related ONTAP commands

- `volume file move show`
- `volume rebalance file-move show`

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Unique identifier of the node.
uuid	string	path	True	Unique identifier of the file move operation.
index	integer	path	True	The index of the file move within the overall operation
fields	array[string]	query	False	Specify the fields to return.

Response

```
Status: 200, Ok
```

Name	Type	Description
cutover_time	integer	Time that the file move operation takes before cutover completes, in seconds.
destination	destination	Destination file information.
elapsed_time	integer	Time elapsed since the start of the file move operation, in seconds.

Name	Type	Description
failure	error	Contains the most recent failure reason for move operation.
files_to_move	files_to_move	A list of source files along with the destination file they are moved to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory. This is only used for FlexVol volume file move operations.
index	integer	An additional unique element identifying one file among many that could possibly be moved as part of a job. File index is an identifier ordered by the file path arrays provided during create. For file move operations that involve only one file, the file-index value of zero is always correct.
is_destination_ready	boolean	Indicates whether the destination file is ready for use.
is_flexgroup	boolean	Indicates whether this is a FlexGroup file move operation.
is_snapshot_fenced	boolean	Indicates whether Snapshot copies are fenced.
max_cutover_time	integer	The maximum amount of time, in seconds that the source can be quiesced before a destination file must be made available for read-write traffic. Not supported in FlexGroup volume file move operations.
max_throughput	integer	Maximum amount of data, in bytes that can be transferred per second in support of this operation. A non-zero value less than 1 MB/s is set to 1 MB/s. A non-zero value greater than 1 MB/s is truncated to the nearest integral megabyte value. If unspecified, the default value is "0" which means no range is set for the data transfer.

Name	Type	Description
node	node	Node hosting the destination of this move operation.
reference	reference	Reference file information.
scanner	scanner	
source	source	Source file information.
svm	svm	SVM of the FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.
uuid	string	The UUID which uniquely identifies the job that started this move operation.
volume	volume	FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.

Example response

```
{
  "cutover_time": 8,
  "destination": {
    "path": "d1/d2/file1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "elapsed_time": 100,
  "failure": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "files_to_move": {
    "destinations": {
      "path": "d1/d2/file1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  }
}
```



```

    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "sources": {
    "path": "d1/d2/file1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
},
"index": 0,
"max_cutover_time": 10,
"max_throughput": 250000,
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"reference": {

```

```

    "max_cutover_time": 5,
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "scanner": {
    "percent": 80,
    "progress": 80000,
    "state": "allocation_map",
    "total": 100000
  },
  "source": {
    "path": "d1/d2/file1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
}

```

```

},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "4fcb3159-a4ee-42b5-bb16-f752f2c430fc",
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

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

volume

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

destination

Destination file information.

Name	Type	Description
path	string	
svm	svm	

Name	Type	Description
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Contains the most recent failure reason for move operation.

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.

destinations

Name	Type	Description
path	string	
svm	svm	
volume	volume	

sources

Name	Type	Description
path	string	
svm	svm	
volume	volume	

files_to_move

A list of source files along with the destination file they are moved to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory. This is only used for FlexVol volume file move operations.

Name	Type	Description
destinations	array[destinations]	The destination file information.
sources	array[sources]	The source file information.

node

Node hosting the destination of this move operation.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

SVM of the source reference file.

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

volume

Volume of the source reference file.

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

reference

Reference file information.

Name	Type	Description
max_cutover_time	integer	The maximum amount of time, in seconds that the source reference file can be quiesced before the corresponding destination file must be made available for read-write traffic. Not supported in FlexGroup volume file move operations.
path	string	The source reference file. If a reference file is specified, data for other files being moved will be transferred as a difference from the reference file. This can save bandwidth and destination storage if the specified source files share blocks. If provided, this input must match one of the source file paths. This input need not be provided if only one source file is specified. Not supported in FlexGroup volume file move operations.
svm	svm	SVM of the source reference file.
volume	volume	Volume of the source reference file.

scanner

Name	Type	Description
percent	integer	Scanner progress, as a percentage.
progress	integer	Scanner progress, in bytes scanned.
state	string	Status of the file move scanner.
total	integer	Total bytes to be scanned.

source

Source file information.

Name	Type	Description
path	string	
svm	svm	
volume	volume	

svm

SVM of the FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.

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

volume

FlexGroup volume in file move operation. Only used in a FlexGroup volume file move operation.

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

error

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

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

Manage FlexCache volumes

Storage flexcache flexcaches endpoint overview

Overview

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

The following relationship configurations are supported:

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

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

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

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

FlexCache APIs

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

- – GET /api/storage/flexcache/flexcaches

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

- – POST /api/storage/flexcache/flexcaches

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

Examples

Creating a FlexCache

The POST request is used to create a FlexCache.

```

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

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

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

```

```

curl -X POST "https://<mgmt-ip>/api/storage/flexcache/flexcaches" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"aggregates\": [ { \"name\": \"aggr_1\" } ], \"name\": \"fc_333\", \"origins\": [ {
\"svm\": { \"name\": \"vs_3\" }, \"volume\": { \"name\": \"vol_o1\" } } ], \"svm\": { \"name\": \"vs_1\" }, \"path\":
\"/fc_333\", \"prepopulate\": { \"dir_paths\": [ \"/dir1\" ] } }" </mgmt-ip>

```

The response:

```

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

```

```

curl -X POST "https://<mgmt-ip>/api/storage/flexcache/flexcaches" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
  \"aggregates\": [ { \"name\": \"aggr_1\" } ], \"name\":
  \"fc_333\", \"origins\": [ { \"svm\": { \"name\": \"vs_3\" },
  \"volume\": { \"name\": \"vol_o1\" } } ], \"svm\": { \"name\": \"vs_1\" },
  \"path\": \"/\
  fc_333\", \"prepopulate\": { \"dir_paths\": [
  \"/dir1\" ], \"exclude_dir_paths\": [ \"/dir1/dir11\" ] } }"

# The response:
{
  \"job\": {
    \"uuid\": \"5afe9ea4-1dcf-11eb-b006-005056ac6a93\",
    \"_links\": {
      \"self\": {
        \"href\": \"/api/cluster/jobs/5afe9ea4-1dcf-11eb-b006-005056ac6a93\"
      }
    }
  }
}

```

```

curl -X POST "https://<mgmt-ip>/api/storage/flexcache/flexcaches" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"aggregates\": [ { \"name\": \"aggr_1\" } ], \"name\": \"fc_333\", \"origins\": [ { \"svm\": { \"name\": \"vs_3\" }, \"volume\": { \"name\": \"vol_o1\" } } ], \"svm\": { \"name\": \"vs_1\" }, \"dr_cache\": true, \"path\": \"/\
fc_333\", \"prepopulate\": { \"dir_paths\": [ \"/dir1\" ] } }" </mgmt-ip>

```

The response:

```

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

```

Retrieving FlexCache attributes

The GET request is used to retrieve FlexCache attributes. The object includes a large set of fields which can be expensive to retrieve. Most notably, the fields `size`, `guarantee.type`, `aggregates`, `path`, `origins.ip_address`, `origins.size`, and `origins.state` are expensive to retrieve. The recommended method to use this API is to filter and retrieve only the required fields.

The API:

`/api/storage/flexcache/flexcaches`

The call:

```

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

```

The response:

```
{ "records": [ { "uuid": "04d5e07b-0ebe-11e9-8180-0050568e0b79", "name": "fc_322", "_links": { "self": { "href": "/api/storage/flexcache/flexcaches/04d5e07b-0ebe-11e9-8180-0050568e0b79" } } }, { "uuid": "47902654-0ea4-11e9-8180-0050568e0b79", "name": "fc_321", "_links": { "self": { "href": "/api/storage/flexcache/flexcaches/47902654-0ea4-11e9-8180-0050568e0b79" } } }, { "uuid": "77e911ff-0ebe-11e9-8180-0050568e0b79", "name": "fc_323", "_links": { "self": { "href": "/api/storage/flexcache/flexcaches/77e911ff-0ebe-11e9-8180-0050568e0b79" } } }, { "uuid": "ddb42bbc-0e95-11e9-8180-0050568e0b79", "name": "fc_32", "_links": { "self": { "href": "/api/storage/flexcache/flexcaches/ddb42bbc-0e95-11e9-8180-0050568e0b79" } } }, { "uuid": "ec774932-0f3c-11e9-8b2b-0050568e0b79", "name": "fc_333", "_links": { "self": { "href": "/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-0050568e0b79" } } } ], "num_records": 5, "_links": { "self": { "href": "/api/storage/flexcache/flexcaches?" } } }
```

Retrieving the attributes of a FlexCache

The GET request is used to retrieve the attributes of a FlexCache. The object includes a large set of fields which can be expensive to retrieve. Most notably, the fields `size`, `guarantee.type`, `aggregates`, `path`, `origins.ip_address`, `origins.size`, and `origins.state` are expensive to retrieve. The recommended method to use this API is to filter and retrieve only the required fields.

The API:

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

The call:

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

The response:

```
{ "uuid": "ec774932-0f3c-11e9-8b2b-0050568e0b79", "name": "fc_333", "svm": { "name": "vs_1", "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79" }, "size": 4294967296, "guarantee": { "type": "volume" }, "dr_cache": "true", "aggregates": [ { "name": "aggr_1", "uuid": "26f34b76-88f8-4a47-b5e0-d8e901fb1114" } ], "origins": [ { "ip_address": "10.140.103.175", "size": 20971520, "create_time": "2019-01-03T15:19:55+05:30", "state": "online", "volume": { "name": "vol_o1", "uuid": "2bc957dd-2617-4afb-8d2f-66ac6070d313" }, "svm": { "name": "vs_3", "uuid": "8aa2cd28-0e92-11e9-b391-0050568e4115" }, "cluster": { "name": "node2", "uuid": "50733f81-0e90-11e9-b391-0050568e4115" } } ], "_links": { "self": { "href": "/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-0050568e0b79" } } }
```

Deleting a FlexCache

The DELETE request is used to delete a FlexCache.

The API:

```
/api/storage/flexcache/flexcaches
```

The call:

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

The response:

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

```
### Modifying a FlexCache volume
Use the PATCH request to update a FlexCache volume.
```

the API:

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

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-0050568e0b79" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"prepopulate\": { \"dir_paths\": [ \"/dir1\" ] } }"</mgmt-ip>
```

The response:

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

```
# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-0050568e0b79"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
 \"prepopulate\": { \"dir_paths\": [ \"/dir1\" ], \"exclude_dir_paths\": [
 \"/dir1/dir11\" ] } }"

# The response:
{
  "job": {
    "uuid": "b574c48c-1da7-11eb-b006-005056ac6a93",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/b574c48c-1da7-11eb-b006-005056ac6a93"
      }
    }
  }
}
```

The call

```
curl -X PATCH "https://<mgmt-ip>/api/storage/flexcache/flexcaches/28f9734a-2fc2-11ed-a5d5-005056bb2b7"
-H "accept: application/json" -H "Content-Type: application/json" -d '{"writeback": {"enabled": true }}'</mgmt-ip>
```

The response:

```
{ "job": { "uuid": "c521a541-3046-11ed-a5d5-005056bbb2b7", "_links": { "self": { "href":
"/api/cluster/jobs/c521a541-3046-11ed-a5d5-005056bbb2b7" } } } }
```

```
# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/28f9734a-2fc2-11ed-a5d5-005056bb2b7"
-H "accept: application/json" -H "Content-Type: application/json" -d '{
  "writeback" : { "enabled" : false } }'
```

```
# The response:
{
  "job": {
    "uuid": "17e193f3-304b-11ed-a5d5-005056bbb2b7",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/17e193f3-304b-11ed-a5d5-005056bbb2b7"
      }
    }
  }
}
```

The call

```
curl -X PATCH "https://<mgmt-ip>/api/storage/flexcache/flexcaches/885dfd0f-ac37-11ed-a2ca-005056bb5573"
-H "accept: application/json" -d '{"relative_size": {"enabled":true, "percentage" : 50}}'</mgmt-ip>
```

The response:

```
{ "job": { "uuid": "34d9c90d-ac38-11ed-a2ca-005056bb5573", "_links": { "self": { "href":
"/api/cluster/jobs/34d9c90d-ac38-11ed-a2ca-005056bb5573" } } } }
```

```

# The call
curl -X PATCH "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/885dfd0f-ac37-11ed-a2ca-005056bb5573"
-H "accept: application/json" -d '{"relative_size" : {"enabled":false}}'

# The response:
{
  "job": {
    "uuid": "14bfcc28-ac82-11ed-83bd-005056bb5573",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/14bfcc28-ac82-11ed-83bd-005056bb5573"
      }
    }
  }
}

```

Retrieve a FlexCache volume in the cluster

GET /storage/flexcache/flexcaches

Introduced In: 9.6

Retrieves FlexCache in the cluster.

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.

- `origins.ip_address` - IP address of origin.
- `origins.size` - Physical size of origin.
- `origins.state` - State of origin.
- `size` - Physical size of FlexCache.
- `guarantee.type` - Space guarantee style of FlexCache.
- `aggregates.name` or `aggregates.uuid` - Name or UUID of aggregate of FlexCache volume.
- `path` - Fully-qualified path of the owning SVM's namespace where the FlexCache is mounted.

Related ONTAP commands

- `volume flexcache show`

Learn more

- [DOC /storage/flexcache/flexcaches](#)

Parameters

Name	Type	In	Required	Description
origins.state	string	query	False	Filter by origins.state
origins.create_time	string	query	False	Filter by origins.create_time
origins.size	integer	query	False	Filter by origins.size
origins.volume.uuid	string	query	False	Filter by origins.volume.uuid
origins.volume.name	string	query	False	Filter by origins.volume.name
origins.ip_address	string	query	False	Filter by origins.ip_address
origins.cluster.name	string	query	False	Filter by origins.cluster.name
origins.cluster.uuid	string	query	False	Filter by origins.cluster.uuid
origins.svm.uuid	string	query	False	Filter by origins.svm.uuid
origins.svm.name	string	query	False	Filter by origins.svm.name
path	string	query	False	Filter by path
uuid	string	query	False	Filter by uuid
guarantee.type	string	query	False	Filter by guarantee.type <ul style="list-style-type: none">• Introduced in: 9.7

Name	Type	In	Required	Description
name	string	query	False	Filter by name <ul style="list-style-type: none"> • maxLength: 203 • minLength: 1
relative_size.percent age	integer	query	False	Filter by relative_size.percent age <ul style="list-style-type: none"> • Introduced in: 9.13
relative_size.enabled	boolean	query	False	Filter by relative_size.enabled <ul style="list-style-type: none"> • Introduced in: 9.13
writeback.enabled	boolean	query	False	Filter by writeback.enabled <ul style="list-style-type: none"> • Introduced in: 9.12
aggregates.name	string	query	False	Filter by aggregates.name
aggregates.uuid	string	query	False	Filter by aggregates.uuid
constituents_per_aggregate	integer	query	False	Filter by constituents_per_aggregate
dr_cache	boolean	query	False	Filter by dr_cache <ul style="list-style-type: none"> • Introduced in: 9.9
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name

Name	Type	In	Required	Description
use_tiered_aggregate	boolean	query	False	Filter by use_tiered_aggregate • Introduced in: 9.8
size	integer	query	False	Filter by size
global_file_locking_enabled	boolean	query	False	Filter by global_file_locking_enabled • Introduced in: 9.9
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
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	return_records	boolean	query	False

Response

Status: 200, Ok

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "aggregates": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "guarantee": {
      "type": "volume"
    },
    "name": "voll",
    "origins": {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "create_time": "2018-06-04 19:00:00 +0000",
      "ip_address": "10.10.10.7",
      "size": 0,
      "state": "error",
      "svm": {
```

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

aggregates

Name	Type	Description
_links	_links	
name	string	
uuid	string	

guarantee

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

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

Name	Type	Description
_links	_links	

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

volume

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

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster management IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	
volume	volume	

prepopulate

FlexCache prepopulate

Name	Type	Description
dir_paths	array[string]	

Name	Type	Description
exclude_dir_paths	array[string]	
recurse	boolean	Specifies whether or not the prepopulate action should search through the <code>dir_paths</code> recursively. If not set, the default value <i>true</i> is used.

relative_size

FlexCache Relative Size

Name	Type	Description
enabled	boolean	Specifies whether the relative sizing is enabled for the FlexCache volume. Relative sizing is introduced as a part of follow the origin feature. When relative sizing is enabled, it blocks any modifications done manually in the absolute size of the FlexCache volume. The size of the FlexCache volume is calculated and entered automatically based on the size of the origin volume.
percentage	integer	Specifies the percent size the FlexCache volume should have relative to the total size of the origin volume. This property is only relevant to a FlexCache volume that has the relative size property enabled.

svm

FlexCache SVM

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

writeback

FlexCache Writeback

Name	Type	Description
enabled	boolean	Indicates whether or not writeback is enabled for the FlexCache volume. Writeback is a storage method where data is first written to the FlexCache volume and then written to the origin of a FlexCache volume.

flexcache

Defines the cache endpoint of FlexCache.

Name	Type	Description
_links	_links	
aggregates	array[aggregates]	
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
dr_cache	boolean	If set to true, a DR cache is created.
global_file_locking_enabled	boolean	Specifies whether or not a FlexCache volume has global file locking mode enabled. Global file locking mode is a mode where protocol read locking semantics are enforced across all FlexCaches and origins of a FlexCache volume. When global file locking mode is enabled, the "is_disconnected_mode_off_for_locks" flag is always set to "true".
guarantee	guarantee	
name	string	FlexCache name
origins	array[flexcache_relationship]	

Name	Type	Description
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
prepopulate	prepopulate	FlexCache prepopulate
relative_size	relative_size	FlexCache Relative Size
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM
use_tiered_aggregate	boolean	Specifies whether or not a Fabricpool-enabled aggregate can be used in FlexCache creation. The <code>use_tiered_aggregate</code> is only used when auto-provisioning a FlexCache volume.
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.
writeback	writeback	FlexCache Writeback

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Create a FlexCache volume in the cluster

POST /storage/flexcache/flexcaches

Introduced In: 9.6

Creates a FlexCache in the cluster.

Required properties

- `name` - Name of FlexCache volume.
- `origins.volume.name` or `origins.volume.uuid` - Name or UUID of origin volume.
- `origins.svm.name` - Name of origin Vserver.
- `svm.name` or `svm.uuid` - Name or UUID of Vserver where FlexCache will be created.

Recommended optional properties

- `path` - Path to mount the FlexCache volume
- `prepopulate.dir_paths` - List of directory-paths to be prepopulated for the FlexCache volume.
- `prepopulate.exclude_dir_paths` - List of directory-paths to be excluded from prepopulation for the FlexCache volume.

Default property values

If not specified in POST, the following default property values are assigned:

- `size` - 10% of origin volume size or 1GB per constituent, whichever is greater.
- `guarantee.type` - none. FlexCache is thin provisioned by default.
- `constituents_per_aggregate` - 4 if `aggregates.name` or `aggregates.uuid` is used.
- `use_tiered_aggregate` - false if `aggr-list` is not used. This property is only used when auto-provisioning a FlexCache volume.
- `is_disconnected_mode_off_for_locks` - false. This property specifies if the origin will honor the cache side locks when doing the lock checks in the disconnected mode.

- `dr_cache` - false if FlexCache is not a DR cache. This property is used to create a DR FlexCache.
- `global_file_locking_enabled` - false. This property specifies whether global file locking is enabled on the FlexCache volume.
- `writeback.enabled` - false. This property specifies whether writeback is enabled for the FlexCache volume.

#* `writeback.per_inode_dirty_limit` - 2500. This property specifies the amount of data in 4KB blocks that the system can write per inode in a FlexCache volume before a writeback is initiated for that inode.

#* `writeback.transfer_limit` - 200. This property specifies the maximum number of 4KB data blocks the system can transfer, at one time, from the cache to the origin. This process will keep on recurring until all the dirty blocks for the inode are transferred to the origin volume.

#* `writeback.scrub_threshold` - 2000000. This property specifies the threshold value in 4KB data blocks which when hit will trigger a scrub that will initiate writeback for all dirty inodes on the FlexCache volume.

- `relative_size.enabled` - false. This property specifies whether the relative sizing is enabled for the FlexCache volume.
- `relative_size.percentage` - 10. This property specifies the percent size FlexCache volume should have relative to the total size of the origin volume.

Related ONTAP commands

- `volume flexcache create`
- `volume flexcache prepopulate start`

Learn more

- [DOC /storage/flexcache/flexcaches](#)

Parameters

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

Request Body

Name	Type	Description
_links	_links	
aggregates	array[aggregates]	

Name	Type	Description
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
dr_cache	boolean	If set to true, a DR cache is created.
global_file_locking_enabled	boolean	Specifies whether or not a FlexCache volume has global file locking mode enabled. Global file locking mode is a mode where protocol read locking semantics are enforced across all FlexCaches and origins of a FlexCache volume. When global file locking mode is enabled, the "is_disconnected_mode_off_for_locks" flag is always set to "true".
guarantee	guarantee	
name	string	FlexCache name
origins	array[flexcache_relationship]	
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
prepopulate	prepopulate	FlexCache prepopulate
relative_size	relative_size	FlexCache Relative Size
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM

Name	Type	Description
use_tiered_aggregate	boolean	Specifies whether or not a Fabricpool-enabled aggregate can be used in FlexCache creation. The use_tiered_aggregate is only used when auto-provisioning a FlexCache volume.
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.
writeback	writeback	FlexCache Writeback

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "guarantee": {
    "type": "volume"
  },
  "name": "vol1",
  "origins": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "create_time": "2018-06-04 19:00:00 +0000",
    "ip_address": "10.10.10.7",
    "size": 0,
    "state": "error",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
```

```

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

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Headers

Name	Description	Type
Location	Useful for tracking the resource location	string

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
66846870	Either the SVM name or origin volume name is missing
66846871	Constituents per aggregate are specified but aggregate name is missing
66846872	More than one origin volume is specified
66846873	The specified SVM UUID is incorrect for the specified SVM name
66846874	The specified aggregate UUID is incorrect for the specified aggregate name
66846875	The specified aggregate name does not exist
66846876	The specified SVM does not exist or is not peered
66846877	The specified origin SVM name is of zero length
66846878	The specified SVM UUID is invalid
66846730	Failed to create a FlexCache volume

Error Code	Description
66846760	The specified SVM is not a data Vserver
66846787	The specified aggregate is a SnapLock aggregate
66846812	The specified aggregate is a Composite aggregate
66846812	The specified junction path is under a FlexCache volume
66846834	FlexCache encryption requires a cluster version of 9.6 or higher
66846835	A volume encryption license is not found

Name	Type	Description
error	error	

Example error

```

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

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregates

Name	Type	Description
_links	_links	
name	string	
uuid	string	

guarantee

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

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

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

volume

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

flexcache_relationship

Name	Type	Description
<code>cluster</code>	cluster	
<code>create_time</code>	string	Creation time of the relationship.
<code>ip_address</code>	string	Cluster management IP of the remote cluster.
<code>size</code>	integer	Size of the remote volume.
<code>state</code>	string	Volume state
<code>svm</code>	svm	
<code>volume</code>	volume	

prepopulate

FlexCache prepopulate

Name	Type	Description
<code>dir_paths</code>	array[string]	
<code>exclude_dir_paths</code>	array[string]	
<code>recurse</code>	boolean	Specifies whether or not the prepopulate action should search through the <code>dir_paths</code> recursively. If not set, the default value <i>true</i> is used.

relative_size

FlexCache Relative Size

Name	Type	Description
enabled	boolean	Specifies whether the relative sizing is enabled for the FlexCache volume. Relative sizing is introduced as a part of follow the origin feature. When relative sizing is enabled, it blocks any modifications done manually in the absolute size of the FlexCache volume. The size of the FlexCache volume is calculated and entered automatically based on the size of the origin volume.
percentage	integer	Specifies the percent size the FlexCache volume should have relative to the total size of the origin volume. This property is only relevant to a FlexCache volume that has the relative size property enabled.

svm

FlexCache SVM

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

writeback

FlexCache Writeback

Name	Type	Description
enabled	boolean	Indicates whether or not writeback is enabled for the FlexCache volume. Writeback is a storage method where data is first written to the FlexCache volume and then written to the origin of a FlexCache volume.

flexcache

Defines the cache endpoint of FlexCache.

Name	Type	Description
_links	_links	
aggregates	array[aggregates]	
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
dr_cache	boolean	If set to true, a DR cache is created.
global_file_locking_enabled	boolean	Specifies whether or not a FlexCache volume has global file locking mode enabled. Global file locking mode is a mode where protocol read locking semantics are enforced across all FlexCaches and origins of a FlexCache volume. When global file locking mode is enabled, the "is_disconnected_mode_off_for_locks" flag is always set to "true".
guarantee	guarantee	
name	string	FlexCache name
origins	array[flexcache_relationship]	

Name	Type	Description
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
prepopulate	prepopulate	FlexCache prepopulate
relative_size	relative_size	FlexCache Relative Size
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM
use_tiered_aggregate	boolean	Specifies whether or not a Fabricpool-enabled aggregate can be used in FlexCache creation. The <code>use_tiered_aggregate</code> is only used when auto-provisioning a FlexCache volume.
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.
writeback	writeback	FlexCache Writeback

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Delete a FlexCache volume

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

Introduced In: 9.6

Deletes a FlexCache. If a FlexCache volume is online, it is offlined before deletion.

Related ONTAP commands

- `volume flexcache delete`

Learn more

- [DOC /storage/flexcache/flexcaches](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Unique identifier of the FlexCache.

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

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default

ONTAP Error Response Codes

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

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve attributes of the FlexCache volume in the cluster

GET /storage/flexcache/flexcaches/{uuid}

Introduced In: 9.6

Retrieves attributes of the FlexCache in the cluster.

Expensive properties

There is an added computational cost to retrieving values for these properties. They are included by default in GET. The recommended method to use this API is to filter and retrieve only the required fields. See [Requesting specific fields](#) to learn more.

- `origins.ip_address` - IP address of origin.
- `origins.size` - Physical size of origin.
- `origins.state` - State of origin.
- `size` - Physical size of FlexCache.
- `guarantee.type` - Space guarantee style of FlexCache.
- `aggregates.name` or `aggregates.uuid` - Name or UUID of aggregate of FlexCache volume.
- `path` - Fully-qualified path of the owning SVM's namespace where the FlexCache is mounted.

Related ONTAP commands

- `volume flexcache show`

Learn more

- [DOC /storage/flexcache/flexcaches](#)

Parameters

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

Response

```
Status: 200, Ok
```

Name	Type	Description
<code>_links</code>	_links	
<code>aggregates</code>	array[aggregates]	

Name	Type	Description
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
dr_cache	boolean	If set to true, a DR cache is created.
global_file_locking_enabled	boolean	Specifies whether or not a FlexCache volume has global file locking mode enabled. Global file locking mode is a mode where protocol read locking semantics are enforced across all FlexCaches and origins of a FlexCache volume. When global file locking mode is enabled, the "is_disconnected_mode_off_for_locks" flag is always set to "true".
guarantee	guarantee	
name	string	FlexCache name
origins	array[flexcache_relationship]	
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
prepopulate	prepopulate	FlexCache prepopulate
relative_size	relative_size	FlexCache Relative Size
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM

Name	Type	Description
use_tiered_aggregate	boolean	Specifies whether or not a Fabricpool-enabled aggregate can be used in FlexCache creation. The use_tiered_aggregate is only used when auto-provisioning a FlexCache volume.
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.
writeback	writeback	FlexCache Writeback

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "guarantee": {
    "type": "volume"
  },
  "name": "vol1",
  "origins": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "create_time": "2018-06-04 19:00:00 +0000",
    "ip_address": "10.10.10.7",
    "size": 0,
    "state": "error",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
```

```

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

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregates

Name	Type	Description
_links	_links	
name	string	
uuid	string	

guarantee

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

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

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

volume

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

flexcache_relationship

Name	Type	Description
<code>cluster</code>	cluster	
<code>create_time</code>	string	Creation time of the relationship.
<code>ip_address</code>	string	Cluster management IP of the remote cluster.
<code>size</code>	integer	Size of the remote volume.
<code>state</code>	string	Volume state
<code>svm</code>	svm	
<code>volume</code>	volume	

prepopulate

FlexCache prepopulate

Name	Type	Description
<code>dir_paths</code>	array[string]	
<code>exclude_dir_paths</code>	array[string]	
<code>recurse</code>	boolean	Specifies whether or not the prepopulate action should search through the <code>dir_paths</code> recursively. If not set, the default value <i>true</i> is used.

relative_size

FlexCache Relative Size

Name	Type	Description
enabled	boolean	Specifies whether the relative sizing is enabled for the FlexCache volume. Relative sizing is introduced as a part of follow the origin feature. When relative sizing is enabled, it blocks any modifications done manually in the absolute size of the FlexCache volume. The size of the FlexCache volume is calculated and entered automatically based on the size of the origin volume.
percentage	integer	Specifies the percent size the FlexCache volume should have relative to the total size of the origin volume. This property is only relevant to a FlexCache volume that has the relative size property enabled.

svm

FlexCache SVM

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

writeback

FlexCache Writeback

Name	Type	Description
enabled	boolean	Indicates whether or not writeback is enabled for the FlexCache volume. Writeback is a storage method where data is first written to the FlexCache volume and then written to the origin of a FlexCache volume.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Pre-populate a FlexCache volume in the cluster

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

Introduced In: 9.8

Prepopulates a FlexCache volume in the cluster, or modifies configuration of the FlexCache volume.

Required properties

- `uuid` - FlexCache volume UUID.

Recommended optional properties

- `prepopulate.exclude_dir_paths` - List of directory-paths to be excluded from prepopulation for the FlexCache volume.
- `prepopulate.dir_paths` - List of directory-paths to be prepopulated for the FlexCache volume.

- `writeback.enabled` - `false`. This property specifies whether writeback is enabled for the FlexCache volume.

#* `writeback.per_inode_dirty_limit` - This property specifies the amount of data in 4KB blocks that the system can write per inode in a FlexCache volume before a writeback is initiated for that inode.

#* `writeback.transfer_limit` - This property specifies the maximum number of 4KB data blocks the system can transfer, at one time, from the cache to the origin. This process will keep on recurring until all the dirty blocks for the inode are transferred to the origin volume.

#* `writeback.scrub_threshold` - This property specifies the threshold value in 4KB data blocks which when hit will trigger a scrub that will initiate writeback for all dirty inodes on the FlexCache volume.

- `relative_size.enabled` - This property specifies whether the relative sizing is enabled for the FlexCache volume.
- `relative_size.percentage` - This property specifies the percent size FlexCache volume should have relative to the total size of the origin volume.

Default property values

If not specified in PATCH, the following default property value is assigned:

- `prepopulate.recurse` - Default value is "true".

Related ONTAP commands

- `volume flexcache prepopulate start`
- `volume flexcache config modify`

Learn more

- [DOC /storage/flexcache/flexcaches](#)

Parameters

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

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

Request Body

Name	Type	Description
_links	_links	
aggregates	array[aggregates]	
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
dr_cache	boolean	If set to true, a DR cache is created.

Name	Type	Description
global_file_locking_enabled	boolean	Specifies whether or not a FlexCache volume has global file locking mode enabled. Global file locking mode is a mode where protocol read locking semantics are enforced across all FlexCaches and origins of a FlexCache volume. When global file locking mode is enabled, the "is_disconnected_mode_off_for_locks" flag is always set to "true".
guarantee	guarantee	
name	string	FlexCache name
origins	array[flexcache_relationship]	
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
prepopulate	prepopulate	FlexCache prepopulate
relative_size	relative_size	FlexCache Relative Size
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM
use_tiered_aggregate	boolean	Specifies whether or not a Fabricpool-enabled aggregate can be used in FlexCache creation. The use_tiered_aggregate is only used when auto-provisioning a FlexCache volume.
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.

Name	Type	Description
writeback	writeback	FlexCache Writeback

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "guarantee": {
    "type": "volume"
  },
  "name": "vol1",
  "origins": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "create_time": "2018-06-04 19:00:00 +0000",
    "ip_address": "10.10.10.7",
    "size": 0,
    "state": "error",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
```

```

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

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
66846922	FlexCache volume does not exist in the SVM
66846923	Failed to prepopulate FlexCache volume because the origin volume is not reachable
66846924	FlexCache volume is offline
66846925	FlexCache volume is not mounted or the junction-path is not active
66846926	The junction-path of FlexCache volume is not active
66846927	FlexCache volume does not have an origin volume
66846928	FlexCache volume does not exist
66846929	Using FlexCache prepopulate requires an effective cluster version of 9.8.0 or later
66846930	Using FlexCache prepopulate in a MetroCluster configuration requires an effective cluster version of 9.8.0 or later on both the local and remote clusters
66846931	Internal Error. FlexCache prepopulate job queue failed. Wait a few minutes, and then try the operation again
66846936	Failed to lookup root file handle for origin of FlexCache volume. Wait a few minutes, and then try the operation again

Error Code	Description
66846937	Internal error. Failed to initialize thread
66846939	Internal error. Failed to get the MSID of the origin volume for FlexCache volume
66846943	Failed to prepopulate because dir_path does not exist
66846944	Failed to get root constituent for FlexCache volume
66846945	Origin of FlexCache volume is not mounted or the junction-path is not active
66846946	The junction-path of origin of FlexCache volume is not active
66846947	FlexCache prepopulate job for FlexCache volume already exists
66846948	FlexCache prepopulate job for FlexCache volume could not be queued because the node is offline

Name	Type	Description
error	error	

Example error

```

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

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregates

Name	Type	Description
_links	_links	
name	string	
uuid	string	

guarantee

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

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

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

volume

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

flexcache_relationship

Name	Type	Description
<code>cluster</code>	cluster	
<code>create_time</code>	string	Creation time of the relationship.
<code>ip_address</code>	string	Cluster management IP of the remote cluster.
<code>size</code>	integer	Size of the remote volume.
<code>state</code>	string	Volume state
<code>svm</code>	svm	
<code>volume</code>	volume	

prepopulate

FlexCache prepopulate

Name	Type	Description
<code>dir_paths</code>	array[string]	
<code>exclude_dir_paths</code>	array[string]	
<code>recurse</code>	boolean	Specifies whether or not the prepopulate action should search through the <code>dir_paths</code> recursively. If not set, the default value <i>true</i> is used.

relative_size

FlexCache Relative Size

Name	Type	Description
enabled	boolean	Specifies whether the relative sizing is enabled for the FlexCache volume. Relative sizing is introduced as a part of follow the origin feature. When relative sizing is enabled, it blocks any modifications done manually in the absolute size of the FlexCache volume. The size of the FlexCache volume is calculated and entered automatically based on the size of the origin volume.
percentage	integer	Specifies the percent size the FlexCache volume should have relative to the total size of the origin volume. This property is only relevant to a FlexCache volume that has the relative size property enabled.

svm

FlexCache SVM

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

writeback

FlexCache Writeback

Name	Type	Description
enabled	boolean	Indicates whether or not writeback is enabled for the FlexCache volume. Writeback is a storage method where data is first written to the FlexCache volume and then written to the origin of a FlexCache volume.

flexcache

Defines the cache endpoint of FlexCache.

Name	Type	Description
_links	_links	
aggregates	array[aggregates]	
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
dr_cache	boolean	If set to true, a DR cache is created.
global_file_locking_enabled	boolean	Specifies whether or not a FlexCache volume has global file locking mode enabled. Global file locking mode is a mode where protocol read locking semantics are enforced across all FlexCaches and origins of a FlexCache volume. When global file locking mode is enabled, the "is_disconnected_mode_off_for_locks" flag is always set to "true".
guarantee	guarantee	
name	string	FlexCache name
origins	array[flexcache_relationship]	

Name	Type	Description
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
prepopulate	prepopulate	FlexCache prepopulate
relative_size	relative_size	FlexCache Relative Size
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM
use_tiered_aggregate	boolean	Specifies whether or not a Fabricpool-enabled aggregate can be used in FlexCache creation. The <code>use_tiered_aggregate</code> is only used when auto-provisioning a FlexCache volume.
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.
writeback	writeback	FlexCache Writeback

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage storage FlexCache origin volumes

Storage flexcache origins endpoint overview

Overview

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

The following relationship configurations are supported:

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

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

FlexCache APIs

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

- – GET `/api/storage/flexcache/origins`
- – GET `/api/storage/flexcache/origins/{uuid}`
- – PATCH `/api/storage/flexcache/origins/{uuid}`

Examples

Retrieving origins of FlexCache attributes

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

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

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

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


Retrieving the attributes of an origin volume

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

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

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

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

```

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

```

```

"volume": {
  "name": "fc_321",
  "uuid": "47902654-0ea4-11e9-8180-0050568e0b79"
},
"svm": {
  "name": "vs_1",
  "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
},
"cluster": {
  "name": "node3",
  "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
}
},
{
  "ip_address": "10.140.103.179",
  "size": 4294967296,
  "create_time": "2019-01-03T00:11:38+05:30",
  "state": "online",
  "volume": {
    "name": "fc_322",
    "uuid": "04d5e07b-0ebe-11e9-8180-0050568e0b79"
  },
  "svm": {
    "name": "vs_1",
    "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
  },
  "cluster": {
    "name": "node3",
    "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
  }
},
{
  "ip_address": "10.140.103.179",
  "size": 4294967296,
  "create_time": "2019-01-03T00:14:52+05:30",
  "state": "online",
  "volume": {
    "name": "fc_323",
    "uuid": "77e911ff-0ebe-11e9-8180-0050568e0b79"
  },
  "svm": {
    "name": "vs_1",
    "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
  },
  "cluster": {
    "name": "node3",

```

```

        "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
    }
}
],
"_links": {
    "self": {
        "href": "/api/storage/flexcache/origins/80fcaee4-0dc2-488b-afb8-86d28a34cda8"
    }
}
}
}

```

Modifying origin options of an origin volume

Use the PATCH request to update options of an origin volume.

```

# the API:
/api/storage/flexcache/origins/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/flexcache/origins/1fbc0ebb-2440-11eb-a86c-005056ac8ca0" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"block_level_invalidation\": \"true\" } "

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

```

Retrieve the origin of a FlexCache in the cluster

GET /storage/flexcache/origins

Introduced In: 9.6

Retrieves origin of FlexCache in the cluster.

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.

- `flexcaches.ip_address` - IP address of FlexCache.
- `flexcaches.size` - Physical size of FlexCache.
- `flexcaches.guarantee.type` - Space guarantee style of FlexCache.
- `flexcaches.state` - State of FlexCache.

Related ONTAP commands

- `volume flexcache origin show-caches`

Learn more

- [DOC /storage/flexcache/origins](#)

Parameters

Name	Type	In	Required	Description
<code>svm.uuid</code>	string	query	False	Filter by <code>svm.uuid</code>
<code>svm.name</code>	string	query	False	Filter by <code>svm.name</code>
<code>block_level_invalidation</code>	boolean	query	False	Filter by <code>block_level_invalidation</code> <ul style="list-style-type: none">• Introduced in: 9.9
<code>flexcaches.state</code>	string	query	False	Filter by <code>flexcaches.state</code>
<code>flexcaches.create_time</code>	string	query	False	Filter by <code>flexcaches.create_time</code>
<code>flexcaches.size</code>	integer	query	False	Filter by <code>flexcaches.size</code>
<code>flexcaches.volume.uuid</code>	string	query	False	Filter by <code>flexcaches.volume.uuid</code>

Name	Type	In	Required	Description
flexcaches.volume.name	string	query	False	Filter by flexcaches.volume.name
flexcaches.ip_addresses	string	query	False	Filter by flexcaches.ip_addresses
flexcaches.cluster.name	string	query	False	Filter by flexcaches.cluster.name
flexcaches.cluster.uuid	string	query	False	Filter by flexcaches.cluster.uuid
flexcaches.svm.uuid	string	query	False	Filter by flexcaches.svm.uuid
flexcaches.svm.name	string	query	False	Filter by flexcaches.svm.name
name	string	query	False	Filter by name <ul style="list-style-type: none"> • maxLength: 203 • minLength: 1
global_file_locking_enabled	boolean	query	False	Filter by global_file_locking_enabled <ul style="list-style-type: none"> • Introduced in: 9.9
uuid	string	query	False	Filter by uuid

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

Response

Status: 200, Ok

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "flexcaches": {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "create_time": "2018-06-04 19:00:00 +0000",
      "ip_address": "10.10.10.7",
      "size": 0,
      "state": "error",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      }
    }
  }
}
```



```

    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
},
"name": "voll1, vol_2",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563512"
}
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

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

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

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

volume

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

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

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster management IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	
volume	volume	

svm

Origin volume SVM

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

flexcache_origin

Defines the origin endpoint of FlexCache.

Name	Type	Description
_links	_links	

Name	Type	Description
block_level_invalidation	boolean	Block level invalidation enables the FlexCache volume to retain blocks that are not changed at the FlexCache volume without having to evict them. This means that the FlexCache volume does not have to again incur the computational cost of fetching blocks over the WAN from the FlexCache volume origin on the next client access. Block level invalidation is a property of the origin volume. Without block level invalidation, any write at the origin volume would evict the whole file at the FlexCache volume, since by default, origin volume does a file level invalidation.
flexcaches	array[flexcache_relationship]	
global_file_locking_enabled	boolean	Specifies whether a global file locking option is enabled for an origin volume of a FlexCache volume.
name	string	Origin volume name
svm	svm	Origin volume SVM
uuid	string	Origin volume UUID. Unique identifier for origin of FlexCache.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

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

Retrieve attributes of the origin of a FlexCache in the cluster

GET /storage/flexcache/origins/{uuid}

Introduced In: 9.6

Retrieves attributes of the origin of a FlexCache in the cluster.

Expensive properties

There is an added computational cost to retrieving values for these properties. They are included by default in GET results. The recommended method to use this API is to filter and retrieve only the required fields. See [Requesting specific fields](#) to learn more.

- `flexcaches.ip_address` - IP address of FlexCache.
- `flexcaches.size` - Physical size of FlexCache.
- `flexcaches.guarantee.type` - Space guarantee style of FlexCache.
- `flexcaches.state` - State of FlexCache.
- `flexcaches.dr_cache` - True if the cache is a DR cache.

Related ONTAP commands

- `volume flexcache origin show-caches`

Learn more

- [DOC /storage/flexcache/origins](#)

Parameters

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

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
block_level_invalidation	boolean	Block level invalidation enables the FlexCache volume to retain blocks that are not changed at the FlexCache volume without having to evict them. This means that the FlexCache volume does not have to again incur the computational cost of fetching blocks over the WAN from the FlexCache volume origin on the next client access. Block level invalidation is a property of the origin volume. Without block level invalidation, any write at the origin volume would evict the whole file at the FlexCache volume, since by default, origin volume does a file level invalidation.
flexcaches	array[flexcache_relationship]	
global_file_locking_enabled	boolean	Specifies whether a global file locking option is enabled for an origin volume of a FlexCache volume.
name	string	Origin volume name
svm	svm	Origin volume SVM
uuid	string	Origin volume UUID. Unique identifier for origin of FlexCache.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "flexcaches": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "create_time": "2018-06-04 19:00:00 +0000",
    "ip_address": "10.10.10.7",
    "size": 0,
    "state": "error",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "name": "voll, vol_2",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  }
}
```



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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

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

volume

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

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster management IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	
volume	volume	

svm

Origin volume SVM

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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

Modify origin options for an origin volume in the cluster

PATCH /storage/flexcache/origins/{uuid}

Introduced In: 9.9

Modifies origin options for a origin volume in the cluster.

Required properties

- `uuid` - Origin volume UUID.
- `block_level_invalidation` - Value for the Block Level Invalidation flag - options {true|false}.

Related ONTAP commands

- `volume flexcache origin config modify`

Learn more

- [DOC /storage/flexcache/origins](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Unique identifier of the origin of the FlexCache volume.

Request Body

Name	Type	Description
<code>_links</code>	_links	

Name	Type	Description
block_level_invalidation	boolean	Block level invalidation enables the FlexCache volume to retain blocks that are not changed at the FlexCache volume without having to evict them. This means that the FlexCache volume does not have to again incur the computational cost of fetching blocks over the WAN from the FlexCache volume origin on the next client access. Block level invalidation is a property of the origin volume. Without block level invalidation, any write at the origin volume would evict the whole file at the FlexCache volume, since by default, origin volume does a file level invalidation.
flexcaches	array[flexcache_relationship]	
global_file_locking_enabled	boolean	Specifies whether a global file locking option is enabled for an origin volume of a FlexCache volume.
name	string	Origin volume name
svm	svm	Origin volume SVM
uuid	string	Origin volume UUID. Unique identifier for origin of FlexCache.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "flexcaches": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "create_time": "2018-06-04 19:00:00 +0000",
    "ip_address": "10.10.10.7",
    "size": 0,
    "state": "error",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "name": "vol1, vol_2",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  }
}
```

```

    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563512"
}

```

Response

Status: 200, Ok

Name	Type	Description
job	job_link	

Example response

```

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

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
66846870	Either origin volume uuid or BLI option value is missing
66847020	Failed to get origin volume details using volume uuid
66847021	Failed to modify origin volume options

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

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

volume

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

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster management IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	
volume	volume	

svm

Origin volume SVM

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

flexcache_origin

Defines the origin endpoint of FlexCache.

Name	Type	Description
_links	_links	

Name	Type	Description
block_level_invalidation	boolean	Block level invalidation enables the FlexCache volume to retain blocks that are not changed at the FlexCache volume without having to evict them. This means that the FlexCache volume does not have to again incur the computational cost of fetching blocks over the WAN from the FlexCache volume origin on the next client access. Block level invalidation is a property of the origin volume. Without block level invalidation, any write at the origin volume would evict the whole file at the FlexCache volume, since by default, origin volume does a file level invalidation.
flexcaches	array[flexcache_relationship]	
global_file_locking_enabled	boolean	Specifies whether a global file locking option is enabled for an origin volume of a FlexCache volume.
name	string	Origin volume name
svm	svm	Origin volume SVM
uuid	string	Origin volume UUID. Unique identifier for origin of FlexCache.

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

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

Manage storage pools

Storage pools endpoint overview

Retrieving storage pool information

The Storage Pools GET API retrieves all shared storage pools in the cluster.

The collection GET returns the storage pool identifiers, UUID and name. The instance GET, by default, returns all of the properties defined in the `storage_pool` object. The "show_spare" query returns a response outside of the records body, which includes the groups of usable spares in the cluster. The usable count for each class of spares does not include reserved spare capacity recommended by ONTAP best practices.

Creating storage pools

Creating a shared storage pool is recommended when distributing flash capacity across the cache tiers of HDD aggregates across an HA pair. POST can be used with specific properties to create a storage pool as requested. At a minimum, the storage pool name, disk count, and the nodes where it should reside, are required to create a new instance.

When using POST with input properties, three properties are required. These are:

- name - Name of the storage pool.
- node.name or node.uuid - Node that can use capacity from the storage pool in their cache tiers.
- capacity.disk_count - Number of disks to be used to create the storage pool.

Examples

Retrieving a list of storage pools from the cluster

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

```

{
"records": [
{
"uuid": "8255fef7-4737-11ec-bd1b-005056bbb879",
"nodes": [
{
"uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
"name": "node-1",
},
{
"uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
"name": "node-2",
}
],
"_links": {
"self": {
"href": "/api/storage/pools/8255fef7-4737-11ec-bd1b-005056bbb879"
}
}
},
],
"num_records": 1,
"_links": {
"self": {
"href": "/api/storage/pools"
}
}
}
}

```

```

# The API:
/api/storage/pools

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

# The response:
{
"records": [
{
"uuid": "8255fef7-4737-11ec-bd1b-005056bbb879",
"name": "new_sp",
"nodes": [
{

```

```

    "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
    "name": "node-1",
  },
  {
    "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
    "name": "node-2",
  }
],
"storage_type": "ssd",
"capacity": {
  "remaining": 1846542336,
  "total": 7386169344,
  "spare_allocation_units": [
    {
      "node": {
        "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
        "name": "node-1",
      },
      "count": 1,
      "syncmirror_pool": "pool0",
      "size": 1846542336,
      "available_size": 1846542336
    },
    {
      "node": {
        "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
        "name": "node-2",
      },
      "count": 0,
      "syncmirror_pool": "pool0",
      "size": 1846542336,
      "available_size": 0
    }
  ],
  "used_allocation_units": [
    {
      "aggregate": {
        "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
        "name": "test_a"
      },
      "allocated_unit_count": 2,
      "node": {
        "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
        "name": "node-1",
      },
      "capacity": 2769813504
    }
  ]
}

```

```

},
{
  "aggregate": {
    "uuid": "f4cc30d5-b052-493a-a49f-19781425f987",
    "name": "test_b"
  },
  "allocated_unit_count": 1,
  "node": {
    "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
    "name": "node-2",
  },
  "capacity": 1384906752
}
],
"disk_count": 4,
"disks": [
  {
    "disk": {
      "name": "VMw-1.11"
    },
    "usable_size": 1902379008,
    "total_size": 1908871168,
  },
  {
    "disk": {
      "name": "VMw-1.12"
    },
    "usable_size": 1902379008,
    "total_size": 1908871168,
  },
  {
    "disk": {
      "name": "VMw-1.23"
    },
    "usable_size": 1902379008,
    "total_size": 1908871168,
  },
  {
    "disk": {
      "name": "VMw-1.24"
    },
    "usable_size": 1902379008,
    "total_size": 1908871168,
  }
]
},

```

```

    "health": {
      "state": "normal",
      "is_healthy": true
    },
  ],
  "num_records": 1,
}

```

Retrieving the usable spare information for the cluster

The following example shows the response from retrieving usable spare information according to ONTAP best practices.

```

# The API:
/api/storage/pools?show_spares=true

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools?show_spares=true" -H
"accept: application/json"

# The response:
{
  "records": [],
  "num_records": 0,
  "spares": [
    {
      "disk_class": "solid_state",
      "disk_type": "ssd",
      "size": 3720609792,
      "checksum_style": "block",
      "syncmirror_pool": "pool0",
      "usable": 12,
      "nodes": [
        {
          "uuid": "54af4069-c1f7-11ec-884e-005056bb6e0b",
          "name": "node-1",
          "_links": {
            "self": {
              "href": "/api/cluster/nodes/54af4069-c1f7-11ec-884e-
005056bb6e0b"
            }
          }
        },
        {
          "uuid": "d50f1acb-c1f6-11ec-9dfd-005056bb8d04",

```



```

        "name": "node-2",
        "_links": {
            "self": {
                "href": "/api/cluster/nodes/d50f1acb-c1f6-11ec-9dfd-
005056bb8d04"
            }
        }
    ]
}
],
"_link": {
    "self": {
        "href": "/api/storage/aggregates?show_spare=true"
    }
}
}

```

Simulating the creation of a storage pool

The following example shows the response containing the simulated layout details of a new storage pool in the cluster.



Each storage pool UUID provided in this response is not guaranteed to be the same UUID for the storage pool if it is created.

```

# The API:
/api/storage/pools

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/pools?simulate=true&fields=*"
-d "{\"nodes\": [{\"name\": \"node1\"}, {\"name\": \"node2\"}], \"name\":
\"storage_pool_1\", \"capacity\": {\"disk_count\": \"4\"}}" -H "accept:
application/json"

# The response:
{
  "records": [
    {
      "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
      "name": "new_sp",
      "nodes": [
        {
          "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
          "name": "node-1",
        },

```

```

    {
      "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
      "name": "node-2",
    }
  ],
  "capacity": {
    "total": 7386169344,
    "disk_count": 4,
    "disks": [
      {
        "disk": {
          "name": "VMw-1.11"
        },
      },
      {
        "disk": {
          "name": "VMw-1.12"
        },
      },
      {
        "disk": {
          "name": "VMw-1.23"
        },
      },
      {
        "disk": {
          "name": "VMw-1.24"
        },
      }
    ]
  }
}
]
}

```

Retrieve storage pools for the entire cluster

GET /storage/pools

Introduced In: 9.11

Retrieves the collection of storage pools for the entire cluster.

Related ONTAP commands

- storage pool show

Parameters

Name	Type	In	Required	Description
uuid	string	query	False	Filter by uuid
name	string	query	False	Filter by name
nodes.uuid	string	query	False	Filter by nodes.uuid
nodes.name	string	query	False	Filter by nodes.name
storage_type	string	query	False	Filter by storage_type
health.is_healthy	boolean	query	False	Filter by health.is_healthy
health.unhealthy_reason.arguments.message	string	query	False	Filter by health.unhealthy_reason.arguments.message
health.unhealthy_reason.arguments.code	string	query	False	Filter by health.unhealthy_reason.arguments.code
health.unhealthy_reason.message	string	query	False	Filter by health.unhealthy_reason.message
health.unhealthy_reason.code	string	query	False	Filter by health.unhealthy_reason.code
health.unhealthy_reason.target	string	query	False	Filter by health.unhealthy_reason.target
health.state	string	query	False	Filter by health.state
capacity.used_allocation_units.node.uuid	string	query	False	Filter by capacity.used_allocation_units.node.uuid

Name	Type	In	Required	Description
capacity.used_allocation_units.node.name	string	query	False	Filter by capacity.used_allocation_units.node.name
capacity.used_allocation_units.aggregate.name	string	query	False	Filter by capacity.used_allocation_units.aggregate.name
capacity.used_allocation_units.aggregate.uuid	string	query	False	Filter by capacity.used_allocation_units.aggregate.uuid
capacity.used_allocation_units.count	integer	query	False	Filter by capacity.used_allocation_units.count
capacity.used_allocation_units.current_usage	integer	query	False	Filter by capacity.used_allocation_units.current_usage
capacity.disk_count	integer	query	False	Filter by capacity.disk_count
capacity.spare_allocation_units.size	integer	query	False	Filter by capacity.spare_allocation_units.size
capacity.spare_allocation_units.count	integer	query	False	Filter by capacity.spare_allocation_units.count
capacity.spare_allocation_units.syncmirror_pool	string	query	False	Filter by capacity.spare_allocation_units.syncmirror_pool
capacity.spare_allocation_units.available_size	integer	query	False	Filter by capacity.spare_allocation_units.available_size

Name	Type	In	Required	Description
capacity.spare_allocation_units.node.uuid	string	query	False	Filter by capacity.spare_allocation_units.node.uuid
capacity.spare_allocation_units.node.name	string	query	False	Filter by capacity.spare_allocation_units.node.name
capacity.disks.total_size	integer	query	False	Filter by capacity.disks.total_size
capacity.disks.usable_size	integer	query	False	Filter by capacity.disks.usable_size
capacity.disks.disk.name	string	query	False	Filter by capacity.disks.disk.name
capacity.total	integer	query	False	Filter by capacity.total
capacity.remaining	integer	query	False	Filter by capacity.remaining
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • 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	
error	error	
num_records	integer	Number of shared storage pools in the cluster.
records	array[storage_pool]	
spares	array[storage_pool_show_spares]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "capacity": {
      "disks": {
        "disk": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "1.0.1"
        },
        "total_size": 0,
        "usable_size": 0
      },
      "remaining": 0,
      "spare_allocation_units": {
        "available_size": 0,
        "node": {
          "_links": {
            "self": {
```

```

        "href": "/api/resourcelink"
    },
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": 0,
"syncmirror_pool": "pool0"
},
"total": 0,
"used_allocation_units": {
    "aggregate": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "current_usage": 0,
    "node": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
}
},
"health": {
    "state": "normal",
    "unhealthy_reason": {
        "arguments": {
            "code": "string",
            "message": "string"
        },
        "code": "4",
        "message": "entry doesn't exist",
        "target": "uuid"
    }
},
"nodes": {
    "_links": {

```



```

    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"storage_type": "SSD",
"uuid": "string"
},
"spares": {
  "checksum_style": "block",
  "disk_class": "solid_state",
  "disk_type": "fc",
  "nodes": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "size": 10156769280,
  "syncmirror_pool": "pool0",
  "usable": 9
}
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
10944513	Unable to retrieve shared disk capability information.
10944514	Unable to enable shared disk capability.
10944527	Storage pools are not supported in MetroCluster configurations.
10944528	Unable to retrieve MetroCluster configuration information.

Error Code	Description
11206662	There is no storage pool matching the specified UUID or name.
11206667	Storage pool feature is not enabled.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

_links

Name	Type	Description
self	href	

disk

Reference to the constituent disk object.

Name	Type	Description
_links	_links	
name	string	

storage_pool_disk

Name	Type	Description
disk	disk	Reference to the constituent disk object.
total_size	integer	Raw capacity of the disk, in bytes.
usable_size	integer	Usable capacity of this disk, in bytes.

node

Specifies what node can use this set of allocation units.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pool_spare_allocation_unit

Name	Type	Description
available_size	integer	The usable capacity of this set of allocation units.
count	integer	The number of spare allocation units on this node.
node	node	Specifies what node can use this set of allocation units.
size	integer	Size of each allocation unit.
syncmirror_pool	string	The RAID SyncMirror Pool to which this allocation unit is assigned.

aggregate

The aggregate that is using this cache capacity.

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	
uuid	string	

node

The node hosting the aggregate using this set of allocation units.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pool_used_allocation_unit

Name	Type	Description
aggregate	aggregate	The aggregate that is using this cache capacity.
count	integer	The number of allocation units used by this aggregate.
current_usage	integer	The amount of cache space used by this aggregate.
node	node	The node hosting the aggregate using this set of allocation units.

capacity

Name	Type	Description
disk_count	integer	The number of disks in the storage pool.
disks	array[storage_pool_disk]	Properties of each disk used in the shared storage pool.
remaining	integer	Remaining usable capacity in the flash pool, in bytes.
spare_allocation_units	array[storage_pool_spare_allocation_unit]	Properties of spare allocation units.

Name	Type	Description
total	integer	Total size of the flash pool, in bytes.
used_allocation_units	array[storage_pool_used_allocation_unit]	Information about the storage pool allocation units participating in the cache tier of an aggregate.

error

Indicates why the storage pool is unhealthy. This property is not returned for healthy storage pools.

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.

health

Properties that outline shared storage pool health.

Name	Type	Description
is_healthy	boolean	Indicates whether the storage pool is able to participate in provisioning operations.
state	string	The state of the shared storage pool.
unhealthy_reason	error	Indicates why the storage pool is unhealthy. This property is not returned for healthy storage pools.

node_reference

Name	Type	Description
_links	_links	
name	string	

Name	Type	Description
uuid	string	

storage_pool

Name	Type	Description
_links	_links	
capacity	capacity	
health	health	Properties that outline shared storage pool health.
name	string	Storage pool name.
nodes	array[node_reference]	Nodes that can use this storage pool for their aggregates.
storage_type	string	Storage type for the disks used to create the storage pool.
uuid	string	Storage pool UUID.

storage_pool_show_spares

Available spares for storage pool.

Name	Type	Description
checksum_style	string	The checksum type that has been assigned to the spares.
disk_class	string	Disk class of spares.
disk_type	string	Type of disk.
nodes	array[node_reference]	Nodes that can use the usable spares for storage pool.
size	integer	Usable size of each spare, in bytes.
syncmirror_pool	string	SyncMirror spare pool.

Name	Type	Description
usable	integer	<p>Total number of usable spares in the bucket. The usable count for each class of spares does not include reserved spare capacity recommended by ONTAP best practices.</p> <ul style="list-style-type: none"> • example: 9 • readOnly: 1 • Introduced in: 9.12 • x-nullable: true

Create a new storage pool

POST /storage/pools

Introduced In: 9.11

Creates a new storage pool using available solid state capacity attached to the nodes specified.

Required properties

The following properties are required in the POST body:

- name - Name of the new storage pool.
- nodes[].name or nodes[].uuid - Nodes that can use cache capacity from the new storage pool. Only nodes in the same HA pair can be specified for a given storage pool. Spare cache capacity will be distributed evenly among the specified nodes.
- capacity.disk_count - Number of SSDs to be used to create the storage pool.

Related ONTAP commands

- storage pool create

Example:

```
POST /api/storage/pools {"nodes": [{"name": "node1"}, {"name": "node2"}],
"name": "storage_pool_1", "capacity": {"disk_count": "4"}}
```

Parameters

Name	Type	In	Required	Description
disk_size	integer	query	False	If set, POST only selects SSDs within five percent of the specified size.
simulate	boolean	query	False	When set to "true", the end point returns a simulated layout of the proposed new storage pool, without changing system state. <ul style="list-style-type: none"> • Introduced in: 9.12
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. <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
_links	_links	
capacity	capacity	
health	health	Properties that outline shared storage pool health.
name	string	Storage pool name.
nodes	array[node_reference]	Nodes that can use this storage pool for their aggregates.
storage_type	string	Storage type for the disks used to create the storage pool.
uuid	string	Storage pool UUID.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "capacity": {
    "disks": {
      "disk": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "1.0.1"
      },
      "total_size": 0,
      "usable_size": 0
    },
    "remaining": 0,
    "spare_allocation_units": {
      "available_size": 0,
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "size": 0,
      "syncmirror_pool": "pool0"
    },
    "total": 0,
    "used_allocation_units": {
      "aggregate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  }
}
```

```

    },
    "current_usage": 0,
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "health": {
    "state": "normal",
    "unhealthy_reason": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    }
  },
  "nodes": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "storage_type": "SSD",
  "uuid": "string"
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "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
11206666	Storage pool is unhealthy.
11208658	A storage pool already uses the specified name.
11208660	Disk does not exist.
11208661	Disk is not a spare disk.
11208662	Disk is not an SSD.
11208663	Disk is reserved for core dump.

Error Code	Description
11208664	Could not determine checksum type of disk.
11208666	Could not determine usable size of disk.
11208668	Could not determine connectivity between controller and disk.
11208670	Could not determine original owner of disk.
11208671	Could not determine SyncMirror pool of disk.
11208673	Could not determine HA mode of node.
11208674	Could not determine HA partner of node.
11208675	Disks specified in the disk list are not visible to node.
11208678	The disk list contains disks from nodes which are not in HA pair.
11208679	Storage pools are not supported on nodes.
11208680	Internal error. Cannot determine configuration for node.
11208681	Node is not online.
11208682	Internal error. Sharing configuration mismatch.
11208684	Unable to share disk.
11208686	Disk cannot be shared.
11208687	Unable to retrieve expected sharing configuration.
11208688	Storage pool create job failed.
11208690	Not all nodes sharing the storage pool view disk as a shared disk.
11208691	Not enough matching spares available.
11208692	A disk list or count is a required parameter for storage pool creation.
11208693	Invalid number of disks specified.
11208698	Internal error. Missing node name.
11208699	Internal error. Missing partner name for node configured for HA.
11208701	Node is a standalone node. Do not specify other nodes with a standalone node.
11208703	Incorrect number of nodes specified. Specify one node or both nodes in an HA pair.
11208704	Specified nodes are not part of HA relationship. Specify one node or both nodes in an HA pair.
11208705	Disk is a data center SSD, which cannot be used in storage pools.

Error Code	Description
11208706	Disk is a SSD-ZNS, which cannot be used in storage pools.
11215756	Missing a required field for POST request.

Name	Type	Description
error	error	

Example error

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

disk

Reference to the constituent disk object.

Name	Type	Description
_links	_links	
name	string	

storage_pool_disk

Name	Type	Description
disk	disk	Reference to the constituent disk object.
total_size	integer	Raw capacity of the disk, in bytes.
usable_size	integer	Usable capacity of this disk, in bytes.

node

Specifies what node can use this set of allocation units.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pool_spare_allocation_unit

Name	Type	Description
available_size	integer	The usable capacity of this set of allocation units.
count	integer	The number of spare allocation units on this node.
node	node	Specifies what node can use this set of allocation units.
size	integer	Size of each allocation unit.
syncmirror_pool	string	The RAID SyncMirror Pool to which this allocation unit is assigned.

aggregate

The aggregate that is using this cache capacity.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

node

The node hosting the aggregate using this set of allocation units.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pool_used_allocation_unit

Name	Type	Description
aggregate	aggregate	The aggregate that is using this cache capacity.
count	integer	The number of allocation units used by this aggregate.

Name	Type	Description
current_usage	integer	The amount of cache space used by this aggregate.
node	node	The node hosting the aggregate using this set of allocation units.

capacity

Name	Type	Description
disk_count	integer	The number of disks in the storage pool.
disks	array[storage_pool_disk]	Properties of each disk used in the shared storage pool.
remaining	integer	Remaining usable capacity in the flash pool, in bytes.
spare_allocation_units	array[storage_pool_spare_allocation_unit]	Properties of spare allocation units.
total	integer	Total size of the flash pool, in bytes.
used_allocation_units	array[storage_pool_used_allocation_unit]	Information about the storage pool allocation units participating in the cache tier of an aggregate.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Indicates why the storage pool is unhealthy. This property is not returned for healthy storage pools.

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

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

health

Properties that outline shared storage pool health.

Name	Type	Description
is_healthy	boolean	Indicates whether the storage pool is able to participate in provisioning operations.
state	string	The state of the shared storage pool.
unhealthy_reason	error	Indicates why the storage pool is unhealthy. This property is not returned for healthy storage pools.

node_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pool

Name	Type	Description
_links	_links	
capacity	capacity	
health	health	Properties that outline shared storage pool health.
name	string	Storage pool name.
nodes	array[node_reference]	Nodes that can use this storage pool for their aggregates.

Name	Type	Description
storage_type	string	Storage type for the disks used to create the storage pool.
uuid	string	Storage pool UUID.

job_link

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

error

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

Manage specific storage pools

Storage pools UUID endpoint overview

Updating storage pools

The PATCH operation is used to modify properties of the storage pool. There are several properties that can be modified on a storage pool. PATCH operations on a storage pool are restricted when another PATCH operation is in progress. The following is a list of properties that can be modified using the PATCH operation including a brief description for each:

- name - Can be updated to rename the storage pool.
- capacity.disk_count - Can be updated to increase the number of disks in a storage pool.
- capacity.spare_allocation_units[].count - Modifying this value requires that the user specify capacity.spare_allocation_units[].node as well. Modifying this value redistributes spare cache capacity among the nodes specified in the operation. When expanding a storage pool, the cache tiers of all aggregates using the storage pool's allocation units are expanded automatically.

Simulated storage pool expansion

The PATCH operation also supports simulated expansion of a storage pool. Running PATCH with the query "simulate" set to "true", and "capacity.disk_count" set to the final disk count will return a response containing the projected new capacity and the new constituent disk list for the storage pool.

Deleting storage pools

If cache capacity from a storage pool is being used in an aggregate, it cannot be deleted. See the /storage/aggregates API for details on deleting aggregates.

Examples

Retrieving a specific pool from the cluster

The following example shows the response of the requested storage pool. If there is no storage pool with the requested UUID, an error is returned.

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

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools/870dd9f2-bdfa-4167-b692-57d1cec874d4" -H "accept: application/json"

# The response:
{
  "uuid": "8255fef7-4737-11ec-bd1b-005056bbb879",
  "name": "new_sp",
  "storage_type": "ssd",
  "nodes": [
    {
      "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
      "name": "node-1",
    },
    {
      "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
      "name": "node-2",
    }
  ],
  "capacity": {
    "remaining": 1846542336,
    "total": 7386169344,
    "spare_allocation_units": [
      {
        "node": {
          "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
```

```

    "name": "node-1",
  },
  "count": 1,
  "syncmirror_pool": "pool0",
  "size": 1846542336,
  "available_size": 1846542336
},
{
  "node": {
    "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
    "name": "node-2",
  },
  "count": 0,
  "syncmirror_pool": "pool0",
  "size": 1846542336,
  "available_size": 0
}
],
"used_allocation_units": [
{
  "aggregate": {
    "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
    "name": "test_a"
  },
  "allocated_unit_count": 2,
  "node": {
    "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
    "name": "node-1",
  },
  "capacity": 2769813504
},
{
  "aggregate": {
    "uuid": "f4cc30d5-b052-493a-a49f-19781425f987",
    "name": "test_b"
  },
  "allocated_unit_count": 1,
  "node": {
    "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
    "name": "node-2",
  },
  "capacity": 1384906752
}
],
"disk_count": 4,
"disks": [

```

```

{
  "disk": {
    "name": "VMw-1.11"
  },
  "usable_size": 1902379008,
  "total_size": 1908871168,
},
{
  "disk": {
    "name": "VMw-1.12"
  },
  "usable_size": 1902379008,
  "total_size": 1908871168,
},
{
  "disk": {
    "name": "VMw-1.23"
  },
  "usable_size": 1902379008,
  "total_size": 1908871168,
},
{
  "disk": {
    "name": "VMw-1.24"
  },
  "usable_size": 1902379008,
  "total_size": 1908871168,
}
]
},
"health": {
  "state": "normal",
  "is_healthy": true
},
}

```

Simulating storage pool expansion

The following example shows the response for a simulated storage pool expansion based on the values of the 'capacity.disk_count' attribute passed in. The query does not modify the existing storage pool, but rather returns how it will look after the expansion. This will be reflected in the following attributes:

- capacity.total- Total space, in bytes.
- capacity.remaining - New remaining capacity, in bytes.
- capacity.disks.disk - New list of constituent disks.
- capacity.disk_count - New number of disks in the pool.

```

# The API:
/api/storage/pools/{uuid}?simulate=true

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/pools/cae60cfe-deae-42bd-
babb-ef437d118314?simulate=true" -H "accept: application/json" -d
"{\"capacity\": {\"disk_count\": 6}}"

# The response:
{
  "records": [
    {
      "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
      "name": "new_sp",
      "capacity": {
        "remaining": 1846542336,
        "total": 7386169344,
        "used_allocation_units": [
          {
            "aggregate": {
              "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
              "name": "test_a"
            },
            "current_usage": 2769813504
          },
          {
            "aggregate": {
              "uuid": "f4cc30d5-b052-493a-a49f-19781425f987",
              "name": "test_b"
            },
            "current_usage": 1384906752
          }
        ],
        "disk_count": 6,
        "disks": [
          {
            "disk": {
              "name": "VMw-1.11"
            }
          },
          {
            "disk": {
              "name": "VMw-1.12"
            }
          }
        ]
      }
    }
  ]
}

```



```
{
  "disk": {
    "name": "VMw-1.23"
  },
},
{
  "disk": {
    "name": "VMw-1.24"
  },
},
{
  "disk": {
    "name": "VMw-1.13"
  },
},
{
  "disk": {
    "name": "VMw-1.14"
  },
}
]
}
```

Adding capacity to a storage pool

The following example shows the workflow of adding disks to the storage pool.

Step 1: Check the current disk count on the storage pool.

```

# The API:
/api/storage/pools

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools/19425837-f2fa-4a9f-8f01-712f626c983c?fields=capacity.disk_count" -H "accept: application/json"

# The response:
{
  "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
  "name": "sp1",
  "capacity": {
    "disk_count": 4
  }
}

```

Step 2: Update the pool with the new disk count in 'capacity.disk_count'. The response to PATCH is a job unless the request is invalid.

```

# The API:
/api/storage/pools

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/pools/19425837-f2fa-4a9f-8f01-712f626c983c" -H "accept: application/hal+json" -d "{\"capacity\":{\"disk_count\": 6}}"

# The response:
{
  "job": {
    "uuid": "c103d15e-730b-11e8-a57f-005056b465d6",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/c103d15e-730b-11e8-a57f-005056b465d6"
      }
    }
  }
}

```

Step 3: Wait for the job to finish, then call GET to see the reflected change.

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

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools/19425837-f2fa-4a9f-8f01-712f626c983c?fields=capacity.disk_count" -H "accept: application/json"

# The response:
{
  "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
  "name": "sp1",
  "capacity": {
    "disk_count": 6
  }
}
```

The following example shows the workflow to redistribute spare capacity among nodes sharing the storage pool Step 1: Check the current spare capacity distribution of the pool.

```

# The API:
/api/storage/pools

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools/f3aafdc6-be35-4d93-9590-5a402bffb4b?fields=capacity.spare_allocation_units" -H "accept: application/json"

# The response:
{
  "uuid": "f3aafdc6-be35-4d93-9590-5a402bffb4b",
  "name": "sp1",
  "capacity": {
    "spare_allocation_units": [
      {
        "node": {
          "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
          "name": "node-1",
        },
        "count": 1,
        "syncmirror_pool": "pool0",
        "size": 1846542336,
        "available_size": 1846542336
      },
      {
        "node": {
          "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
          "name": "node-2",
        },
        "count": 0,
        "syncmirror_pool": "pool0",
        "size": 1846542336,
        "available_size": 0
      }
    ],
  }
}

```

Step 2: Update the pool so that the spare allocation unit count is symmetrically modified for each node. The response to PATCH is a job unless the request is invalid.

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/pools/f3aafdc6-be35-4d93-9590-5a402bfffbe4b" -H "accept: application/hal+json" -d '{ "capacity" : { "spare_allocation_units": [{"node": {"name": "node-1"}, "count":0}, {"node": {"name": "node-2"}, "count": 1}]}}'
```

```
# The response:
{
  "job": {
    "uuid": "6b7ab28e-168d-11ea-8a50-0050568eca76",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/6b7ab28e-168d-11ea-8a50-0050568eca76"
      }
    }
  }
}
```

Step 3: Wait for the job to finish, then call GET to see the reflected change.

```

# The API:
/api/storage/pools

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/pools/f3aafdc6-be35-4d93-9590-5a402bffb4b?fields=capacity.spare_allocation_units" -H "accept: application/json"

# The response:
{
  "uuid": "f3aafdc6-be35-4d93-9590-5a402bffb4b",
  "name": "sp1",
  "capacity": {
    "spare_allocation_units": [
      {
        "node": {
          "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
          "name": "node-1",
        },
        "count": 0,
        "syncmirror_pool": "pool0",
        "size": 1846542336,
        "available_size": 0
      },
      {
        "node": {
          "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
          "name": "node-2",
        },
        "count": 1,
        "syncmirror_pool": "pool0",
        "size": 1846542336,
        "available_size": 1846542336
      }
    ],
  }
}

```

Delete a storage pool specified by the UUID

DELETE /storage/pools/{uuid}

Introduced In: 9.11

Deletes the storage pool specified by the UUID. This request starts a job and returns a link to that job.

Related ONTAP commands

- `storage pool delete`

Parameters

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

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
11209668	Capacity from pool is being used by one or more aggregates.
11209670	Unable to stop sharing a constituent disk of the storage pool.
11209671	Cannot find a node sharing this storage pool.
11209672	Disk is not shared.
11209673	Unable to assign shared capacity from a constituent disk.
11209675	A constituent disk of the storage pool has non-spare capacity.
11209676	A constituent disk of the storage pool is reserved for coredump.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a storage pool specified by the UUID

GET /storage/pools/{uuid}

Introduced In: 9.11

Retrieves the storage pool specified by the UUID.

Related ONTAP commands

- `storage pool show -storage-pool-uuid`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Storage pool UUID.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
capacity	capacity	
health	health	Properties that outline shared storage pool health.
name	string	Storage pool name.
nodes	array[node_reference]	Nodes that can use this storage pool for their aggregates.
storage_type	string	Storage type for the disks used to create the storage pool.
uuid	string	Storage pool UUID.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "capacity": {
    "disks": {
      "disk": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "1.0.1"
      },
      "total_size": 0,
      "usable_size": 0
    },
    "remaining": 0,
    "spare_allocation_units": {
      "available_size": 0,
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "size": 0,
      "syncmirror_pool": "pool0"
    },
    "total": 0,
    "used_allocation_units": {
      "aggregate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  }
}
```

```

    },
    "current_usage": 0,
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "health": {
    "state": "normal",
    "unhealthy_reason": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    }
  },
  "nodes": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "storage_type": "SSD",
  "uuid": "string"
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
11206662	There is no storage pool matching the specified UUID or name.
11215856	The specified storage pool is not healthy.

Name	Type	Description
error	error	

Example error

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

disk

Reference to the constituent disk object.

Name	Type	Description
_links	_links	
name	string	

storage_pool_disk

Name	Type	Description
disk	disk	Reference to the constituent disk object.
total_size	integer	Raw capacity of the disk, in bytes.
usable_size	integer	Usable capacity of this disk, in bytes.

node

Specifies what node can use this set of allocation units.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pool_spare_allocation_unit

Name	Type	Description
available_size	integer	The usable capacity of this set of allocation units.
count	integer	The number of spare allocation units on this node.
node	node	Specifies what node can use this set of allocation units.
size	integer	Size of each allocation unit.
syncmirror_pool	string	The RAID SyncMirror Pool to which this allocation unit is assigned.

aggregate

The aggregate that is using this cache capacity.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

node

The node hosting the aggregate using this set of allocation units.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pool_used_allocation_unit

Name	Type	Description
aggregate	aggregate	The aggregate that is using this cache capacity.
count	integer	The number of allocation units used by this aggregate.

Name	Type	Description
current_usage	integer	The amount of cache space used by this aggregate.
node	node	The node hosting the aggregate using this set of allocation units.

capacity

Name	Type	Description
disk_count	integer	The number of disks in the storage pool.
disks	array[storage_pool_disk]	Properties of each disk used in the shared storage pool.
remaining	integer	Remaining usable capacity in the flash pool, in bytes.
spare_allocation_units	array[storage_pool_spare_allocation_unit]	Properties of spare allocation units.
total	integer	Total size of the flash pool, in bytes.
used_allocation_units	array[storage_pool_used_allocation_unit]	Information about the storage pool allocation units participating in the cache tier of an aggregate.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Indicates why the storage pool is unhealthy. This property is not returned for healthy storage pools.

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

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

health

Properties that outline shared storage pool health.

Name	Type	Description
is_healthy	boolean	Indicates whether the storage pool is able to participate in provisioning operations.
state	string	The state of the shared storage pool.
unhealthy_reason	error	Indicates why the storage pool is unhealthy. This property is not returned for healthy storage pools.

node_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

error

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

Update a storage pool specified by the UUID

PATCH /storage/pools/{uuid}

Introduced In: 9.11

Updates the storage pool specified by the UUID with the properties in the body. This request starts a job and returns a link to that job.

Related ONTAP commands

- storage pool rename
- storage pool reassign
- storage pool add

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Storage pool UUID.
simulate	boolean	query	False	When set to "true" in conjunction with a PATCH on "capacity.disk_count", the end point returns a simulated layout of the storage pool with the additional disks, without changing system state. • Introduced in: 9.12

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

Request Body

Name	Type	Description
_links	_links	
capacity	capacity	
health	health	Properties that outline shared storage pool health.
name	string	Storage pool name.
nodes	array[node_reference]	Nodes that can use this storage pool for their aggregates.

Name	Type	Description
storage_type	string	Storage type for the disks used to create the storage pool.
uuid	string	Storage pool UUID.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "capacity": {
    "disks": {
      "disk": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "1.0.1"
      },
      "total_size": 0,
      "usable_size": 0
    },
    "remaining": 0,
    "spare_allocation_units": {
      "available_size": 0,
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "size": 0,
      "syncmirror_pool": "pool0"
    },
    "total": 0,
    "used_allocation_units": {
      "aggregate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  }
}
```

```
    },
    "current_usage": 0,
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "health": {
    "state": "normal",
    "unhealthy_reason": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    }
  },
  "nodes": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "storage_type": "SSD",
  "uuid": "string"
}
```

Response

Status: 200, Ok

Name	Type	Description
job	job_link	
num_records	integer	Number of shared storage pools in the cluster.
records	array[storage_pool]	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "capacity": {
      "disks": {
        "disk": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "1.0.1"
        },
        "total_size": 0,
        "usable_size": 0
      },
      "remaining": 0,
      "spare_allocation_units": {
        "available_size": 0,
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "size": 0,
        "syncmirror_pool": "pool0"
      }
    },
  },
}
```

```

"total": 0,
"used_allocation_units": {
  "aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "current_usage": 0,
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"health": {
  "state": "normal",
  "unhealthy_reason": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
},
"nodes": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"storage_type": "SSD",
"uuid": "string"

```

```
}  
}
```

Response

Status: 202, Accepted

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
11211658	Node does not have enough spare capacity.
11211659	Valid allocation unit input is required.
11211662	Specified node is not part of the storage pool.
11211663	Failed to reassign available capacity in the storage pool.
11211664	Could not fix the broken allocation unit for the storage pool.
11212673	Could not grow one or more aggregates.
11212679	Adding specified number of disks will expand storage pool beyond maximum supported disk limit.
11212680	Incorrect node specified.
11212681	0 is an invalid value for disk_count.
11212682	Adding the specified number of disks will result in the storage pool reaching the maximum disk limit reserved for RAID-TEC use only. At this limit, the storage pool can only allocate capacity to aggregates containing RAID-TEC RAID groups. Existing aggregates containing RAID groups other than RAID-TEC will not automatically grow to the new capacity.
11212683	Renaming storage pool to new name failed.
11212763	Storage pool add job failed.
11215657	Storage pool PATCH request have missing parameters.
11215658	Storage pool PATCH request for reassign is invalid.

Error Code	Description
11215659	Storage pool PATCH request for reassign have invalid allocation unit count.
11215660	Storage pool PATCH request for reassign have invalid node name.
11215662	Storage pool PATCH request have invalid disk count.

Name	Type	Description
error	error	

Example error

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

disk

Reference to the constituent disk object.

Name	Type	Description
_links	_links	
name	string	

storage_pool_disk

Name	Type	Description
disk	disk	Reference to the constituent disk object.
total_size	integer	Raw capacity of the disk, in bytes.
usable_size	integer	Usable capacity of this disk, in bytes.

node

Specifies what node can use this set of allocation units.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pool_spare_allocation_unit

Name	Type	Description
available_size	integer	The usable capacity of this set of allocation units.
count	integer	The number of spare allocation units on this node.
node	node	Specifies what node can use this set of allocation units.
size	integer	Size of each allocation unit.
syncmirror_pool	string	The RAID SyncMirror Pool to which this allocation unit is assigned.

aggregate

The aggregate that is using this cache capacity.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

node

The node hosting the aggregate using this set of allocation units.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pool_used_allocation_unit

Name	Type	Description
aggregate	aggregate	The aggregate that is using this cache capacity.
count	integer	The number of allocation units used by this aggregate.

Name	Type	Description
current_usage	integer	The amount of cache space used by this aggregate.
node	node	The node hosting the aggregate using this set of allocation units.

capacity

Name	Type	Description
disk_count	integer	The number of disks in the storage pool.
disks	array[storage_pool_disk]	Properties of each disk used in the shared storage pool.
remaining	integer	Remaining usable capacity in the flash pool, in bytes.
spare_allocation_units	array[storage_pool_spare_allocation_unit]	Properties of spare allocation units.
total	integer	Total size of the flash pool, in bytes.
used_allocation_units	array[storage_pool_used_allocation_unit]	Information about the storage pool allocation units participating in the cache tier of an aggregate.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Indicates why the storage pool is unhealthy. This property is not returned for healthy storage pools.

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

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

health

Properties that outline shared storage pool health.

Name	Type	Description
is_healthy	boolean	Indicates whether the storage pool is able to participate in provisioning operations.
state	string	The state of the shared storage pool.
unhealthy_reason	error	Indicates why the storage pool is unhealthy. This property is not returned for healthy storage pools.

node_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pool

Name	Type	Description
_links	_links	
capacity	capacity	
health	health	Properties that outline shared storage pool health.
name	string	Storage pool name.
nodes	array[node_reference]	Nodes that can use this storage pool for their aggregates.

Name	Type	Description
storage_type	string	Storage type for the disks used to create the storage pool.
uuid	string	Storage pool UUID.

job_link

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

error

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

Manage storage ports

Storage ports endpoint overview


Retrieving storage port information

The storage port GET API retrieves all of the storage ports in the cluster.

Examples

1) Retrieve a list of storage ports from the cluster

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



```

# The API:
/api/storage/ports

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

# The response:
{
"records": [
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
      }
    },
    "name": "0a",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0a"
      }
    }
  },
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
        }
      }
    },
    "name": "0b",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0b"
      }
    }
  }
]
}

```

```

    }
  },
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
        }
      }
    },
    "name": "0c",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0c"
      }
    }
  },
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
        }
      }
    },
    "name": "0d",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0d"
      }
    }
  },
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {

```

```

    "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
  }
}
},
"name": "0e",
"_links": {
  "self": {
    "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0e"
  }
}
},
{
  "node": {
    "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
    "name": "node-1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
      }
    }
  },
  "name": "0f",
  "_links": {
    "self": {
      "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-
00a0985a72ee/0f"
    }
  }
},
{
  "node": {
    "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
    "name": "node-1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-
00a0985a72ee"
      }
    }
  },
  "name": "0g",
  "_links": {
    "self": {

```

```
    "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0g"
  }
}
},
],
"num_records": 7,
"_links": {
  "self": {
    "href": "/api/storage/ports"
  }
}
}
```

2) Retrieve a specific storage port from the cluster

The following example shows the response of the requested storage port. If there is no storage port with the requested node uuid and name, an error is returned.

```
# The API:
/api/storage/ports/{node.uuid}/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0a" -H "accept: application/hal+json"

# The response:
{
  "node": {
    "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
    "name": "node-1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
      }
    }
  },
  "name": "0a",
  "description": "SAS Host Adapter 0a (PMC-Sierra PM8001 rev. C)",
  "wwn": "500a098003633df0",
  "speed": 6,
  "cable": {
    "part_number": "112-00429+A0",
    "serial_number": "629230774",
    "identifier": "500a0980066e2c01-500a098003633df0",
    "length": "0.5m"
  },
  "state": "online",
  "enabled": true,
  "firmware_version": "01.12.09.00",
  "type": "sas",
  "redundant": true,
  "in_use": true,
  "_links": {
    "self": {
      "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0a"
    }
  }
}
```

Updating a storage port

The storage port PATCH API modifies the port mode for storage/network use and allows the port to be enabled/disabled.

Examples

1) Using an Ethernet port for storage

The following example sets an Ethernet port mode for storage use:

```
# The API:
/api/storage/ports/{node.uuid}/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/e3a" -H "accept: application/hal+json" -H "Content-Type: application/hal_json" -d '{"mode": "storage"}'

# The response:
{
}
```

2) Disabling a storage port

The following example disables an unused storage port:

```
# The API:
/api/storage/ports/{node.uuid}/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/e3a" -H "accept: application/hal+json" -H "Content-Type: application/hal_json" -d '{"enabled": "false"}'

# The response:
{
}
```

Retrieve storage ports

GET /storage/ports

Introduced In: 9.6

Retrieves a collection of storage ports.

Related ONTAP commands

- `storage port show`

Learn more

- [DOC /storage/ports](#)

Parameters

Name	Type	In	Required	Description
speed	number	query	False	Filter by speed
in_use	boolean	query	False	Filter by in_use <ul style="list-style-type: none">• Introduced in: 9.9
serial_number	string	query	False	Filter by serial_number
description	string	query	False	Filter by description
part_number	string	query	False	Filter by part_number
mode	string	query	False	Filter by mode <ul style="list-style-type: none">• Introduced in: 9.8
mac_address	string	query	False	Filter by mac_address
cable.identifier	string	query	False	Filter by cable.identifier
cable.length	string	query	False	Filter by cable.length
cable.part_number	string	query	False	Filter by cable.part_number

Name	Type	In	Required	Description
cable.serial_number	string	query	False	Filter by cable.serial_number
error.message	string	query	False	Filter by error.message
error.corrective_action	string	query	False	Filter by error.corrective_action
board_name	string	query	False	Filter by board_name
type	string	query	False	Filter by type <ul style="list-style-type: none"> • Introduced in: 9.9
node.uuid	string	query	False	Filter by node.uuid
node.name	string	query	False	Filter by node.name
enabled	boolean	query	False	Filter by enabled <ul style="list-style-type: none"> • Introduced in: 9.9
firmware_version	string	query	False	Filter by firmware_version <ul style="list-style-type: none"> • Introduced in: 9.9
name	string	query	False	Filter by name
state	string	query	False	Filter by state
redundant	boolean	query	False	Filter by redundant <ul style="list-style-type: none"> • Introduced in: 9.9
wwpn	string	query	False	Filter by wwpn <ul style="list-style-type: none"> • Introduced in: 9.9

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

Response

Status: 200, Ok

Name	Type	Description
_links	_links	

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "board_name": "string",
    "cable": {
      "identifier": "500a0980000b6c3f-50000d1703544b80",
      "length": "2m",
      "part_number": "112-00431+A0",
      "serial_number": 616930439
    },
    "description": "SAS Host Adapter 2a (PMC-Sierra PM8072 rev. C)",
    "firmware_version": "03.08.09.00",
    "mac_address": "string",
    "mode": "storage",
    "name": "2a",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "part_number": "111-03801",
    "serial_number": "7A2463CC45B",
    "speed": 6,
    "state": "online",
    "type": "sas",
    "wwn": "50000d1703544b80",
    "wwpn": "string"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

error

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

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_port

Name	Type	Description
board_name	string	
cable	cable	
description	string	
enabled	boolean	
error	error	
firmware_version	string	
force	boolean	
in_use	boolean	Specifies whether any devices are connected through this port
mac_address	string	
mode	string	Operational mode of a non-dedicated Ethernet port
name	string	
node	node	
part_number	string	
redundant	boolean	Specifies whether all devices connected through this port have a redundant path from another port
serial_number	string	
speed	number	Operational port speed in Gbps
state	string	
type	string	
wwn	string	World Wide Name
wwpn	string	World Wide Port Name

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a storage port

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

Introduced In: 9.6

Retrieves a specific storage port.

Related ONTAP commands

- `storage port show`

Learn more

- [DOC /storage/ports](#)

Parameters

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

Response

Status: 200, Ok

Name	Type	Description
board_name	string	
cable	cable	

Name	Type	Description
description	string	
enabled	boolean	
error	error	
firmware_version	string	
force	boolean	
in_use	boolean	Specifies whether any devices are connected through this port
mac_address	string	
mode	string	Operational mode of a non-dedicated Ethernet port
name	string	
node	node	
part_number	string	
redundant	boolean	Specifies whether all devices connected through this port have a redundant path from another port
serial_number	string	
speed	number	Operational port speed in Gbps
state	string	
type	string	
wwn	string	World Wide Name
wwpn	string	World Wide Port Name

Example response

```
{
  "board_name": "string",
  "cable": {
    "identifier": "500a0980000b6c3f-50000d1703544b80",
    "length": "2m",
    "part_number": "112-00431+A0",
    "serial_number": 616930439
  },
  "description": "SAS Host Adapter 2a (PMC-Sierra PM8072 rev. C)",
  "firmware_version": "03.08.09.00",
  "mac_address": "string",
  "mode": "storage",
  "name": "2a",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "part_number": "111-03801",
  "serial_number": "7A2463CC45B",
  "speed": 6,
  "state": "online",
  "type": "sas",
  "wwn": "50000d1703544b80",
  "wwpn": "string"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

error

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

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Update a storage port

PATCH /storage/ports/{node.uuid}/{name}

Introduced In: 9.11

Updates a storage port.

Related ONTAP commands

- `storage port modify`
- `storage port enable`
- `storage port disable`

Learn more

- [DOC /storage/ports](#)

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Node UUID
name	string	path	True	Port name

Request Body

Name	Type	Description
board_name	string	
cable	cable	
description	string	
enabled	boolean	
error	error	

Name	Type	Description
firmware_version	string	
force	boolean	
in_use	boolean	Specifies whether any devices are connected through this port
mac_address	string	
mode	string	Operational mode of a non-dedicated Ethernet port
name	string	
node	node	
part_number	string	
redundant	boolean	Specifies whether all devices connected through this port have a redundant path from another port
serial_number	string	
speed	number	Operational port speed in Gbps
state	string	
type	string	
wwn	string	World Wide Name
wwpn	string	World Wide Port Name

Example request

```
{
  "board_name": "string",
  "cable": {
    "identifier": "500a098000b6c3f-50000d1703544b80",
    "length": "2m",
    "part_number": "112-00431+A0",
    "serial_number": 616930439
  },
  "description": "SAS Host Adapter 2a (PMC-Sierra PM8072 rev. C)",
  "firmware_version": "03.08.09.00",
  "mac_address": "string",
  "mode": "storage",
  "name": "2a",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "part_number": "111-03801",
  "serial_number": "7A2463CC45B",
  "speed": 6,
  "state": "online",
  "type": "sas",
  "wwn": "50000d1703544b80",
  "wwpn": "string"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
17891328	Port operation "<operation>\\" failed on port \\"<name>\". This might indicate a hardware error, an illegal request, or an aborted command.</name></operation>
17891329	Port "<name>\\" is invalid.</name>
17891330	Port operation "<operation>\\" failed on port \\"<port>\\" because it is not supported on this port type.</port></operation>
17891331	Cannot complete operation on port "<port>\\". The status of the port is not available.</port>
17891332	Port operation "<operation>\\" failed on port \\"<port>\\". The port is not offline.</port></operation>
17891333	Port operation "<operation>\\" failed on port \\"<port>\\". The port is already offline.</port></operation>
17891334	Port operation "<operation>\\" failed on port \\"<port>\\". One or more devices on the port is currently being sanitized.</port></operation>
17891338	Device operation "<operation>\\" failed on port \\"<name>\\". Some devices can only be accessed through this port. Disabling this port might cause multiple device failures.</name></operation>
17891339	Device operation "<operation>\\" failed on port \\"<name>\\". Some devices can only be accessed through this port. Disabling this port might cause multiple device failures. Use the \\"-force\" parameter to ignore checks and <operation>the port.</operation></name></operation>
17891344	Device operation "<operation>\\" failed on port \\"<name>\\" because the port is in use.</name></operation>
17891345	Device operation "<operation>\\" failed on port \\"<name>\\" because the port is in use. Use the \\"-force\" parameter to ignore checks and <operation>the port.</operation></name></operation>
17891346	Port "<name>\\" is already enabled.</name>
17891347	Port "<name>\\" is already disabled.</name>
17891350	Port "<name>\\" is not a valid storage Ethernet port.</name>
17891350	Port "<port>\\" is not a valid storage Ethernet port.</port>

Error Code	Description
17891352	Port operation "<operation>" failed on port \"<name>\". Reboot node \"<node>\" to recover.</node></name></operation>
17891353	Unable to enable port "<port>" because it is in network mode.</port>
17891354	Unable to disable port "<port>" because it is in network mode.</port>
17891355	Port operation "<operation>" failed on port \"<name>\" because it is not supported on dedicated ports.</name></operation>
17891356	Unable to <operation>port \"<port>\" when setting mode to \"<mode>\".</mode></port></operation>

Name	Type	Description
error	error	

Example error

```

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

```

Definitions

See Definitions

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

error

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

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_port

Name	Type	Description
board_name	string	
cable	cable	
description	string	
enabled	boolean	

Name	Type	Description
error	error	
firmware_version	string	
force	boolean	
in_use	boolean	Specifies whether any devices are connected through this port
mac_address	string	
mode	string	Operational mode of a non-dedicated Ethernet port
name	string	
node	node	
part_number	string	
redundant	boolean	Specifies whether all devices connected through this port have a redundant path from another port
serial_number	string	
speed	number	Operational port speed in Gbps
state	string	
type	string	
wwn	string	World Wide Name
wwpn	string	World Wide Port Name

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

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

Manage QoS policies

Storage QoS policies endpoint overview

Quality of Service Configuration

A QoS policy defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy is associated. There are two types of policies that can be configured: fixed, which defines a fixed SLO, or adaptive which defines a variable SLO for a storage object. Adaptive policies vary the SLO depending on the space usage of the storage object. A policy can be either a fixed policy or an adaptive one, not both. Service level objectives include minimum and maximum limits on throughput in terms of IOPS. Only maximum limits can be set in terms of both IOPS and/or throughput (MB/s). A QoS policy can be used to enforce SLOs for multiple storage objects by specifying "capacity_shared" to true. For example, if a QoS policy with "capacity_shared" is set to true and it has maximum_throughput_iops set to 1000, and this policy is assigned to four volumes, then the combined throughput of all four volumes is limited to 1000 IOPS. If "capacity_shared" is set to false then, each storage object will have its SLOs enforced individually. For example, in the previous case if the same policy was applied to four volumes but with "capacity_shared" set to false, then each of the volumes would be limited to 1000 IOPS individually. Once "capacity_shared" is set, it cannot be modified. Adaptive parameters can specify the variable SLOs in terms of IOPS/TB. The actual IOPS enforced on the storage object can be calculated using the allocated space on the storage object. The policies are enforced individually amongst storage objects.

Examples

1) Create a fixed QoS policy

The following example shows how to create a fixed QoS policy to limit throughput for a storage object between 5000 IOPS and 10000 IOPS which has capacity_shared set to false. This QoS policy can be used as a template to apply on multiple storage objects to provide individual SLOs to each object.

```
curl -X POST
"https://172.21.69.245/api/storage/qos/policies?return_timeout=0" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"fixed\": { \"capacity_shared\": false, \"max_throughput_iops\": 10000,
\"min_throughput_iops\": 5000 }, \"name\":
\"qos_policy_5000_to_10000_iops\", \"svm\": { \"name\": \"vs0\" }}"
```

2) Create an adaptive QoS policy

The following example shows how to create an adaptive QoS policy which provides 5000 IOPS per GB of allocated space for a storage object with a peak of 6000 IOPS. Minimum IOPS regardless of allocated space are 1000 IOPS.

```
curl -X POST
"https://172.21.69.245/api/storage/qos/policies?return_timeout=0" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"adaptive\": { \"absolute_min_iops\": 1000, \"expected_iops\": 5000,
\"expected_iops_allocation\": \"used_space\", \"peak_iops\": 6000,
\"peak_iops_allocation\": \"allocated_space\" }, \"name\":
\"adaptive_pg_5k_to_6k\", \"svm\": { \"name\": \"vs0\" }}"
```

-

3) Update an existing QoS policy

The following example shows how to update SLOs of an existing QoS policy and also rename it.

```
curl -X PATCH "https://172.21.69.245/api/storage/qos/policies/d38bafc0-
5a51-11e9-bd5b-005056ac6f1f?return_timeout=0" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"fixed\": {
\"max_throughput_iops\": 15000, \"min_throughput_iops\": 10000 },
\"name\": \"qos_policy_10k_to_15k_iops\"}"
```

4) Delete an existing QoS policy

When a QoS policy is deleted any associations of the policy with a storage objects are also removed.

```
curl -X DELETE "https://172.21.69.245/api/storage/qos/policies/d38bafc0-
5a51-11e9-bd5b-005056ac6f1f?return_timeout=0" -H "accept:
application/json"
```

Retrieve QoS policies

GET /storage/qos/policies

Introduced In: 9.6

Retrieves a collection of QoS policies.

Parameters

Name	Type	In	Required	Description
name	string	query	False	Filter by name
uuid	string	query	False	Filter by uuid
fixed.max_throughput_mbps	integer	query	False	Filter by fixed.max_throughput_mbps
fixed.min_throughput_iops	integer	query	False	Filter by fixed.min_throughput_iops
fixed.min_throughput_mbps	integer	query	False	Filter by fixed.min_throughput_mbps <ul style="list-style-type: none">• Introduced in: 9.8
fixed.capacity_shared	boolean	query	False	Filter by fixed.capacity_shared
fixed.max_throughput_iops	integer	query	False	Filter by fixed.max_throughput_iops
scope	string	query	False	Filter by scope <ul style="list-style-type: none">• Introduced in: 9.11
policy_class	string	query	False	Filter by policy_class <ul style="list-style-type: none">• Introduced in: 9.10

Name	Type	In	Required	Description
object_count	integer	query	False	Filter by object_count
adaptive.expected_iops	integer	query	False	Filter by adaptive.expected_iops
adaptive.peak_iops_allocation	string	query	False	Filter by adaptive.peak_iops_allocation • Introduced in: 9.10
adaptive.absolute_min_iops	integer	query	False	Filter by adaptive.absolute_min_iops
adaptive.peak_iops	integer	query	False	Filter by adaptive.peak_iops
adaptive.block_size	string	query	False	Filter by adaptive.block_size • Introduced in: 9.10
adaptive.expected_iops_allocation	string	query	False	Filter by adaptive.expected_iops_allocation • Introduced in: 9.10
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
pgid	integer	query	False	Filter by pgid • Introduced in: 9.10
fields	array[string]	query	False	Specify the fields to return.

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

Response

Status: 200, Ok

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "adaptive": {
      "block_size": "any",
      "expected_iops_allocation": "used_space",
      "peak_iops_allocation": "used_space"
    },
    "name": "extreme",
    "object_count": 0,
    "pgid": 0,
    "policy_class": "undefined",
    "scope": "cluster",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
}
```

```
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

_links

Name	Type	Description
self	href	

adaptive

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

Name	Type	Description
absolute_min_iops	integer	Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
block_size	string	Specifies the block size
expected_iops	integer	Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
expected_iops_allocation	string	Specifies the size to be used to calculate expected IOPS per TB. The size options are either the storage object allocated space or the storage object used space.
peak_iops	integer	Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size.
peak_iops_allocation	string	Specifies the size to be used to calculate peak IOPS per TB. The size options are either the storage object allocated space or the storage object used space.

fixed

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

Name	Type	Description
capacity_shared	boolean	Specifies whether the capacities are shared across all objects that use this QoS policy-group. Default is false.

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

svm

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

qos_policy

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

Name	Type	Description
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.
pgid	integer	Policy group ID of the QoS policy.
policy_class	string	Class of the QoS policy.
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.
svm	svm	
uuid	string	

Create a QoS policy

POST /storage/qos/policies

Introduced In: 9.6

Creates a QoS policy.

Required properties

- `svm.uuid` or `svm.name` - The existing SVM owning the QoS policy.
- `name` - The name of the QoS policy.
- `fixed.*` or `adaptive.*` - Either of the fixed or adaptive parameters.

Default property values

- If `fixed.*` parameters are specified, then `capacity.shared` is set to `false` by default.

Related ONTAP commands

- `qos policy-group create`
- `qos adaptive-policy-group create`

Parameters

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

Request Body

Name	Type	Description
_links	_links	

Name	Type	Description
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.
pgid	integer	Policy group ID of the QoS policy.
policy_class	string	Class of the QoS policy.
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.
svm	svm	
uuid	string	

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "adaptive": {
    "block_size": "any",
    "expected_iops_allocation": "used_space",
    "peak_iops_allocation": "used_space"
  },
  "name": "extreme",
  "object_count": 0,
  "pgid": 0,
  "policy_class": "undefined",
  "scope": "cluster",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Headers

Name	Description	Type
Location	Useful for tracking the resource location	string

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
8454147	The maximum limit for QoS policies has been reached.
8454154	The name specified for creating conflicts with an existing QoS policy name.
8454260	Invalid value for maximum and minimum fields. Valid values for max_throughput_iops and max_throughput_mbps combination is for the ratio of max_throughput_mbps and max_throughput_iops to be within 1 to 4096.
8454273	Invalid value for an adaptive field. Value should be non-zero.
8454277	The name specified for creating an adaptive QoS policy conflicts with an existing fixed QoS policy name.
8454278	The name specified for creating a fixed QoS policy conflicts with an existing adaptive QoS policy name.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

adaptive

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

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

Name	Type	Description
peak_iops_allocation	string	Specifies the size to be used to calculate peak IOPS per TB. The size options are either the storage object allocated space or the storage object used space.

fixed

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

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

svm

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

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

qos_policy

Name	Type	Description
_links	_links	
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.
pgid	integer	Policy group ID of the QoS policy.
policy_class	string	Class of the QoS policy.
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.
svm	svm	
uuid	string	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Delete a QoS policy

DELETE /storage/qos/policies/{uuid}

Introduced In: 9.6

Deletes a QoS policy. All QoS workloads associated with the policy are removed.

Related ONTAP commands

- qos policy-group delete
- qos adaptive-policy-group delete

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	<ul style="list-style-type: none"> • Introduced in: 9.8

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

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a QoS policy

GET /storage/qos/policies/{uuid}

Introduced In: 9.6

Retrieves a specific QoS policy.

Related ONTAP commands

- `qos policy-group show`
- `qos adaptive-policy-group show`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	<ul style="list-style-type: none">• Introduced in: 9.8
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

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

Name	Type	Description
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.
svm	svm	
uuid	string	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "adaptive": {
    "block_size": "any",
    "expected_iops_allocation": "used_space",
    "peak_iops_allocation": "used_space"
  },
  "name": "extreme",
  "object_count": 0,
  "pgid": 0,
  "policy_class": "undefined",
  "scope": "cluster",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

adaptive

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

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

Name	Type	Description
peak_iops_allocation	string	Specifies the size to be used to calculate peak IOPS per TB. The size options are either the storage object allocated space or the storage object used space.

fixed

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

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

svm

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

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Update a QoS policy

PATCH /storage/qos/policies/{uuid}

Introduced In: 9.6

Update a specific QoS policy.

Related ONTAP commands

- `qos policy-group modify`
- `qos adaptive-policy-group modify`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	<ul style="list-style-type: none"> • Introduced in: 9.8

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

Request Body

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

Name	Type	Description
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.
pgid	integer	Policy group ID of the QoS policy.
policy_class	string	Class of the QoS policy.
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.
svm	svm	
uuid	string	

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "adaptive": {
    "block_size": "any",
    "expected_iops_allocation": "used_space",
    "peak_iops_allocation": "used_space"
  },
  "name": "extreme",
  "object_count": 0,
  "pgid": 0,
  "policy_class": "undefined",
  "scope": "cluster",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
8454147	The maximum limit for QoS policies has been reached.
8454154	The name specified for creating conflicts with an existing QoS policy name.
8454260	Invalid value for maximum and minimum fields. Valid values for max_throughput_iops and max_throughput_mbps combination is for the ratio of max_throughput_mbps and max_throughput_iops to be within 1 to 4096.
8454273	Invalid value for an adaptive field. Value should be non-zero.
8454277	The name specified for creating an adaptive QoS policy conflicts with an existing fixed QoS policy name.
8454278	The name specified for creating a fixed QoS policy conflicts with an existing adaptive QoS policy name.
8454286	Modifications on these cluster scoped preset policies is prohibited.
8454327	The existing fixed QoS policy cannot be modified to an adaptive QoS policy.
8454328	The existing adaptive QoS policy cannot be modified to a fixed QoS policy.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

adaptive

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

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

Name	Type	Description
peak_iops_allocation	string	Specifies the size to be used to calculate peak IOPS per TB. The size options are either the storage object allocated space or the storage object used space.

fixed

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

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

svm

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

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

qos_policy

Name	Type	Description
_links	_links	
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.
pgid	integer	Policy group ID of the QoS policy.
policy_class	string	Class of the QoS policy.
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.
svm	svm	
uuid	string	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage QoS workloads

Storage QoS workloads endpoint overview

Quality of Service Workloads

A QoS workload represents a storage object that is tracked by QoS.

Examples

Retrieving a list of QoS workloads from the cluster

The following example retrieves all the workloads in the cluster.

```
curl -X GET "https://<mgmt-ip>/api/storage/qos/workloads" -H "accept: application/hal+json"
```

Retrieving a specific QoS workload from the cluster

The following example retrieves a requested workload from the cluster.

```
curl -X GET "https://<mgmt-ip>/api/storage/qos/workloads/77b68b1c-a458-11eb-baaa-005056bb873e" -H "accept: application/hal+json"
```

Retrieve QoS workloads

GET /storage/qos/workloads

Introduced In: 9.10

Retrieves a collection of QoS workloads.

Parameters

Name	Type	In	Required	Description
volume	string	query	False	Filter by volume
lun	string	query	False	Filter by lun
file	string	query	False	Filter by file
uuid	string	query	False	Filter by uuid
name	string	query	False	Filter by name
workload_class	string	query	False	Filter by workload_class
policy.uuid	string	query	False	Filter by policy.uuid
policy.name	string	query	False	Filter by policy.name
wid	integer	query	False	Filter by wid
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
qtree	string	query	False	Filter by qtree
fields	array[string]	query	False	Specify the fields to return.

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

Response

Status: 200, Ok

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "file": "string",
    "lun": "string",
    "name": "volumel-wid123",
    "policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "performance",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "qtree": "string",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  },
}
```

```
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": "volume1",
  "wid": 123,
  "workload_class": "autovolume"
}
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

_links

Name	Type	Description
self	href	

policy

QoS policy group reference.

Name	Type	Description
_links	_links	

Name	Type	Description
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.

svm

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

qos_workload

Name	Type	Description
_links	_links	
file	string	Name of the file.
lun	string	Name of the LUN. The name of the LUN will be displayed as "(unknown)" if the name cannot be retrieved.
name	string	Name of the QoS workload.
policy	policy	QoS policy group reference.
qtree	string	Name of the Qtree.
svm	svm	
uuid	string	
volume	string	Name of the volume. The name of the volume will be displayed as "(unknown)" if the name cannot be retrieved.

Name	Type	Description
wid	integer	Workload ID of the QoS workload.
workload_class	string	Class of the QoS workload.

Retrieve a QoS workload

GET /storage/qos/workloads/{uuid}

Introduced In: 9.10

Retrieves a specific QoS workload.

Related ONTAP command

- qos workload show

Parameters

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

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
file	string	Name of the file.
lun	string	Name of the LUN. The name of the LUN will be displayed as "(unknown)" if the name cannot be retrieved.
name	string	Name of the QoS workload.
policy	policy	QoS policy group reference.

Name	Type	Description
qtree	string	Name of the Qtree.
svm	svm	
uuid	string	
volume	string	Name of the volume. The name of the volume will be displayed as "(unknown)" if the name cannot be retrieved.
wid	integer	Workload ID of the QoS workload.
workload_class	string	Class of the QoS workload.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "file": "string",
  "lun": "string",
  "name": "volume1-wid123",
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "qtree": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": "volume1",
  "wid": 123,
  "workload_class": "autovolume"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

policy

QoS policy group reference.

Name	Type	Description
_links	_links	
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.

svm

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage storage qtrees

Storage qtrees endpoint overview

Overview

A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a FlexVol volume or a FlexGroup volume.

Qtree QoS policy

Qtree QoS policy and settings enforce Service Level Objectives (SLOs) on a qtree. SLOs can be set by specifying "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops" and/or "qos_policy.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_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a qtree is not supported if its containing volume or SVM has a QoS policy attached, or a file or LUN in its containing volume already has a QoS policy attached.

Qtree APIs

The following APIs are used to create, retrieve, modify, and delete qtrees.

– POST /api/storage/qtrees

– GET /api/storage/qtrees

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

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

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

Examples

Creating a qtree inside a volume for an SVM

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

The following example shows how to create a qtree in a FlexVol volume with a given security style, user,

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

# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/qtrees?return_records=true' -H
'accept: application/hal+json' -d @test_qtree_post.txt
test_qtree_post.txt (body):
{
  "svm": {
    "name": "svm1"
  },
  "volume": {
    "name": "fv"
  },
  "name": "qt1",
  "security_style": "unix",
  "user": {
    "name": "unix_user1"
  },
  "group": {
    "name": "unix_group1"
  },
  "unix_permissions": 744,
  "export_policy": {
    "name": "default"
  },
  "qos_policy": {
    "max_throughput_iops": 1000
  }
}

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

```

"name": "qt1",
"security_style": "unix",
"user": {
  "name": "unix_user1"
},
"group": {
  "name": "unix_group1"
},
"unix_permissions": 744,
"export_policy": {
  "name": "default"
},
"qos_policy": {
  "min_throughput_iops": 0,
  "min_throughput_mbps": 0,
  "max_throughput_iops": 1000,
  "max_throughput_mbps": 0,
  "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
  "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
},
"_links": {
  "self": {
    "href": "/api/storage/qtrees/?volume.name=fv&name=qt1"
  }
}
],
"job": {
  "uuid": "84edef3c-4f6d-11e9-9a71-005056a7f717",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/84edef3c-4f6d-11e9-9a71-005056a7f717"
    }
  }
}
}

```

Retrieving qtrees

This API is used to retrieve qtrees.

The following example shows how to retrieve qtrees belonging to SVM *svm1* and volume *fv*. The `svm.name` and `volume.name` query parameters are used to find the required qtrees.

```

# The API:
GET /api/storage/qtrees

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/qtrees/?svm.name=svml&volume.name=fv" -H 'accept:
application/hal+json'

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

```



```

        "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
        }
    },
    "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
            "self": {
                "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
            }
        }
    },
    "id": 1,
    "name": "qt1",
    "_links": {
        "self": {
            "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-
005056a7f717/1"
        }
    }
},
{
    "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
            "self": {
                "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
            }
        }
    },
    "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
            "self": {
                "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
            }
        }
    },
    "id": 2,
    "name": "qt2",

```

```

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

```

Retrieving properties of a specific qtree using a qtree identifier

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

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

```

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

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

```

```
    "self": {
      "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
    }
  },
  "id": 2,
  "name": "qt2",
  "security_style": "unix",
  "user": {
    "name": "unix_user1"
  },
  "group": {
    "name": "unix_group1"
  },
  "unix_permissions": 744,
  "export_policy": {
    "name": "default",
    "id": 12884901889,
    "_links": {
      "self": {
        "href": "/api/protocols/nfs/export-policies/12884901889"
      }
    }
  },
  "qos_policy": {
    "min_throughput_iops": 0,
    "min_throughput_mbps": 0,
    "max_throughput_iops": 1000,
    "max_throughput_mbps": 0,
    "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
    "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52",
    "_links": {
      "self": {
        "href": "/api/storage/qos/policies/39ac471f-ff35-11e9-b0f9-005056a7ab52"
      }
    }
  },
  "statistics": {
    "timestamp": "2019-04-09T05:50:42Z",
    "status": "ok",
    "iops_raw": {
      "read": 0,
      "write": 0,
      "other": 3,
```

```

    "total": 3
  },
  "throughput_raw": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  }
},
"path": "/fv/qt2",
"nas": {
  "path": "/fv/qt2",
},
"_links": {
  "self": {
    "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/2"
  }
}
}
}

```

Retrieving properties of a specific qtree using the qtree name

This API is used to retrieve properties of a specific qtree using "qtree.name". The following example shows how to retrieve all of the properties belonging to qtree "qt2". The `svm.name` and `volume.name` query parameters are used here along with the qtree name.

```

# The API:
GET /api/storage/qtrees/

# The call:
curl -X GET 'https://<mgmt-
ip>/api/storage/qtrees/?svm.name=svml&volume.name=fv&name=qt2&fields=*' -H
'accept: application/hal+json'
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  }
},

```

```
"volume": {
  "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
  "name": "fv",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
    }
  }
},
"id": 2,
"name": "qt2",
"security_style": "unix",
"user": {
  "name": "unix_user1"
},
"group": {
  "name": "unix_group1"
},
"unix_permissions": 744,
"export_policy": {
  "name": "default",
  "id": 12884901889,
  "_links": {
    "self": {
      "href": "/api/protocols/nfs/export-policies/12884901889"
    }
  }
},
"qos_policy": {
  "min_throughput_iops": 0,
  "min_throughput_mbps": 0,
  "max_throughput_iops": 1000,
  "max_throughput_mbps": 0,
  "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
  "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52",
  "_links": {
    "self": {
      "href": "/api/storage/qos/policies/39ac471f-ff35-11e9-b0f9-005056a7ab52"
    }
  }
},
"statistics": {
  "timestamp": "2019-04-09T05:50:42Z",
  "status": "ok",
```

```
"iops_raw": {
  "read": 0,
  "write": 0,
  "other": 3,
  "total": 3
},
"throughput_raw": {
  "read": 0,
  "write": 0,
  "other": 0,
  "total": 0
}
},
"_links": {
  "self": {
    "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/2"
  }
}
}
```

Updating a qtree

This API is used to update a qtree.

The following example shows how to update properties in a qtree.

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

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/2' -H 'accept: application/hal+json' -d '@test_qtree_patch.txt'
test_qtree_patch.txt (body):
{
  "security_style": "mixed",
  "user": {
    "name": "unix_user1"
  },
  "group": {
    "name": "unix_group1"
  },
  "unix_permissions": 777,
  "export_policy": {
    "id": "9",
    "name": "exp1"
  },
  "qos_policy": {
    "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab53"
  }
}
```

Renaming a qtree

This API is used to rename a qtree.

The following example below shows how to rename a qtree with a new name.

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

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/1' -H 'accept: application/hal+json' -d '{ "name": "new_qt1" }'
```

Deleting a qtree inside a volume of an SVM

This API is used to delete a qtree inside a volume of an SVM.

The following example shows how to delete a qtree.

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

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/2" -H 'accept: application/hal+json'
```

Retrieve qtrees

GET /storage/qtrees

Introduced In: 9.6

Retrieves qtrees configured for all FlexVol volumes or FlexGroup volumes.

Use the `fields` query parameter to retrieve all properties of the qtree. If the `fields` query parameter is not used, then GET returns the qtree name and qtree id only.

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.

- `statistics.*`

Related ONTAP commands

- `qtree show`

Parameters

Name	Type	In	Required	Description
<code>security_style</code>	string	query	False	Filter by <code>security_style</code>
<code>name</code>	string	query	False	Filter by name
<code>path</code>	string	query	False	Filter by path

Name	Type	In	Required	Description
qos_policy.name	string	query	False	Filter by qos_policy.name • Introduced in: 9.8
qos_policy.min_throughput_mbps	integer	query	False	Filter by qos_policy.min_throughput_mbps • Introduced in: 9.8
qos_policy.min_throughput_iops	integer	query	False	Filter by qos_policy.min_throughput_iops • Introduced in: 9.8
qos_policy.max_throughput_mbps	integer	query	False	Filter by qos_policy.max_throughput_mbps • Introduced in: 9.8
qos_policy.uuid	string	query	False	Filter by qos_policy.uuid • Introduced in: 9.8
qos_policy.max_throughput_iops	integer	query	False	Filter by qos_policy.max_throughput_iops • Introduced in: 9.8
group.id	string	query	False	Filter by group.id • Introduced in: 9.9
group.name	string	query	False	Filter by group.name • Introduced in: 9.9

Name	Type	In	Required	Description
unix_permissions	integer	query	False	Filter by unix_permissions
volume.uuid	string	query	False	Filter by volume.uuid
volume.name	string	query	False	Filter by volume.name
filesystem_path	string	query	False	Filter by filesystem_path • Introduced in: 9.10
statistics.iops_raw.read	integer	query	False	Filter by statistics.iops_raw.read • Introduced in: 9.8
statistics.iops_raw.other	integer	query	False	Filter by statistics.iops_raw.other • Introduced in: 9.8
statistics.iops_raw.write	integer	query	False	Filter by statistics.iops_raw.write • Introduced in: 9.8
statistics.iops_raw.total	integer	query	False	Filter by statistics.iops_raw.total • Introduced in: 9.8
statistics.timestamp	string	query	False	Filter by statistics.timestamp • Introduced in: 9.8

Name	Type	In	Required	Description
statistics.status	string	query	False	Filter by statistics.status • Introduced in: 9.8
statistics.throughput_raw.read	integer	query	False	Filter by statistics.throughput_raw.read • Introduced in: 9.8
statistics.throughput_raw.other	integer	query	False	Filter by statistics.throughput_raw.other • Introduced in: 9.8
statistics.throughput_raw.write	integer	query	False	Filter by statistics.throughput_raw.write • Introduced in: 9.8
statistics.throughput_raw.total	integer	query	False	Filter by statistics.throughput_raw.total • Introduced in: 9.8
export_policy.name	string	query	False	Filter by export_policy.name
export_policy.id	integer	query	False	Filter by export_policy.id
user.id	string	query	False	Filter by user.id • Introduced in: 9.9
user.name	string	query	False	Filter by user.name • Introduced in: 9.9

Name	Type	In	Required	Description
nas.path	string	query	False	Filter by nas.path <ul style="list-style-type: none"> • Introduced in: 9.9
id	integer	query	False	Filter by id <ul style="list-style-type: none"> • Max value: 4994 • Min value: 0
_tags	string	query	False	Filter by _tags <ul style="list-style-type: none"> • Introduced in: 9.13
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Max value: 120 • Min value: 0 • Default value: 1
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[qtree]	

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"
    ],
    "export_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 100,
      "name": "default"
    },
    "filesystem_path": "/dir1/mtree1",
    "group": {
      "id": "20001",
      "name": "unix_group1"
    },
    "id": 1,
    "nas": {
      "path": "/volume3/mtree1"
    },
    "path": "/volume3/mtree1",
    "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"
  },
  "security_style": "unix",
  "statistics": {
    "iops_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_permissions": 493,
  "user": {
    "id": "10001",
    "name": "unix_user1"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}

```

```
}  
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
918235	A volume with UUID was not found.
2621462	The specified SVM does not exist.
5242889	Failed to get the qtree from volume.
5242956	Failed to obtain qtree.
5242965	Invalid qtree path. The volume name component of the qtree path, must be the same as the volume specified with the parameter.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

group

The user set as owner of the qtree.

Name	Type	Description
id	string	The numeric ID of the group that owns the qtree. Valid in POST or PATCH.
name	string	Alphanumeric group name of group that owns the qtree. Valid in POST or PATCH.

nas

Name	Type	Description
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.

qos_policy

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

iops_raw

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

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

throughput_raw

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

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

statistics

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

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

svm

Required in POST

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

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

user

The user set as owner of the qtree.

Name	Type	Description
id	string	The numeric ID of the user who owns the qtree. Valid in POST or PATCH.
name	string	Alphanumeric username of user who owns the qtree. Valid in POST or PATCH.

volume

Required in POST

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

qtree

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

Name	Type	Description
_links	_links	

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.
export_policy	export_policy	Export Policy
filesystem_path	string	Path of the qtree directory. This path is relative to the volume root directory.
group	group	The user set as owner of the qtree.
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
nas	nas	
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH. This field is to be deprecated and replaced with nas.path.
qos_policy	qos_policy	
security_style	string	Security style. Valid in POST or PATCH.
statistics	statistics	These are raw IOPS and throughput performance numbers. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.

Name	Type	Description
user	user	The user set as owner of the qtree.
volume	volume	Required in POST

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Create a qtree in a FlexVol or FlexGroup volume

POST /storage/qtrees

Introduced In: 9.6

Creates a qtree in a FlexVol volume or a FlexGroup volume.

After a qtree is created, the new qtree is assigned an identifier. This identifier is obtained using a qtree GET request. This identifier is used in the API path for the qtree PATCH and DELETE operations.

Required properties

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

Recommended optional properties

If not specified in POST, the values are inherited from the volume.

- `security_style` - Security style for the qtree.
- `unix_permissions` - UNIX permissions for the qtree.
- `export_policy.name` or `export_policy.id` - Export policy of the SVM for the qtree.

Related ONTAP commands

- `qtree create`

Parameters

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

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
_links	_links	
_tags	array[string]	Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings.
export_policy	export_policy	Export Policy
filesystem_path	string	Path of the qtree directory. This path is relative to the volume root directory.
group	group	The user set as owner of the qtree.
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
nas	nas	
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH. This field is to be deprecated and replaced with nas.path.
qos_policy	qos_policy	
security_style	string	Security style. Valid in POST or PATCH.

Name	Type	Description
statistics	statistics	These are raw IOPS and throughput performance numbers. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
user	user	The user set as owner of the qtree.
volume	volume	Required in POST

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "filesystem_path": "/dir1/qtreen1",
  "group": {
    "id": "20001",
    "name": "unix_group1"
  },
  "id": 1,
  "nas": {
    "path": "/volume3/qtreen1"
  },
  "path": "/volume3/qtreen1",
  "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"
  },
  "security_style": "unix",
  "statistics": {
```

```
"iops_raw": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"status": "ok",
"throughput_raw": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 +0000"
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"unix_permissions": 493,
"user": {
  "id": "10001",
  "name": "unix_user1"
},
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```

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

```

Headers

Name	Description	Type
Location	Useful for tracking the resource location	string

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
262245	Invalid field value.
262247	Invalid field value.
917525	The specified volume does not exist in Vserver.
917927	The specified volume was not found.
918232	Either <code>volume.name</code> or <code>volume.uuid</code> must be provided.
918236	The specified <code>volume.uuid</code> and <code>volume.name</code> refer to different volumes.
1703954	Export Policy name specified is invalid.
2621462	The specified SVM does not exist.

Error Code	Description
2621706	The specified <code>svm.uuid</code> and <code>svm.name</code> do not refer to the same SVM.
2621707	No SVM was specified. Either <code>svm.name</code> or <code>svm.uuid</code> must be provided.
5242881	Cannot create qtree because the volume is read-only.
5242886	Failed to create qtree.
5242894	Qtree with empty name "" is not allowed, as that is reserved for the default qtree.
5242900	Qtree not supported on FlexCache volume
5242948	Qtree is not supported on FlexCache origin volume.
5242951	Export Policy supplied does not belong to the specified Export Policy ID.
5242952	Export Policy ID specified is invalid.
5242953	Qtree name must be provided.
5242967	UNIX user or group ID must be 32-bit unsigned integer.
5242970	FlexCache create is in progress for the volume.
6622064	Security-style ntfs is not supported on SnapMirror Business Continuity (SMBC) relationship volume.
8454348	QoS on qtrees is not supported because not all nodes in the cluster can support it.
9437324	The security style unified is not supported.
23724050	Failed to resolve user or group name.
66846755	Failed to determine whether volume is a FlexCache volume or not.
66846839	Failed to determine the effective cluster version of all the nodes hosting FlexCache volumes connected to FlexCache origin volume.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

group

The user set as owner of the qtree.

Name	Type	Description
id	string	The numeric ID of the group that owns the qtree. Valid in POST or PATCH.
name	string	Alphanumeric group name of group that owns the qtree. Valid in POST or PATCH.

nas

Name	Type	Description
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.

qos_policy

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>max_throughput_iops</code>	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
<code>max_throughput_mbps</code>	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
<code>min_throughput_iops</code>	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.
<code>min_throughput_mbps</code>	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
<code>name</code>	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
<code>uuid</code>	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

`iops_raw`

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

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

throughput_raw

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

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

statistics

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

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

svm

Required in POST

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

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

user

The user set as owner of the qtree.

Name	Type	Description
id	string	The numeric ID of the user who owns the qtree. Valid in POST or PATCH.
name	string	Alphanumeric username of user who owns the qtree. Valid in POST or PATCH.

volume

Required in POST

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

qtree

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

Name	Type	Description
_links	_links	

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.
export_policy	export_policy	Export Policy
filesystem_path	string	Path of the qtree directory. This path is relative to the volume root directory.
group	group	The user set as owner of the qtree.
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
nas	nas	
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH. This field is to be deprecated and replaced with nas.path.
qos_policy	qos_policy	
security_style	string	Security style. Valid in POST or PATCH.
statistics	statistics	These are raw IOPS and throughput performance numbers. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.

Name	Type	Description
user	user	The user set as owner of the qtree.
volume	volume	Required in POST

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Delete a qtree

DELETE /storage/qtrees/{volume.uuid}/{id}

Introduced In: 9.6

Deletes a qtree.

Related ONTAP commands

- `qtree delete`

Parameters

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

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```

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

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
917505	Vserver not found.
917506	Volume not found.
917525	The specified volume does not exist in Vserver.
918235	A volume with UUID was not found.
5242894	The default qtree cannot be deleted.
5242895	Failed to delete the qtree.
5242897	This operation is not permitted on read-only volume.
5242898	This operation is only permitted on a data Vserver.
5242916	Cannot delete qtree because the volume contains one or more LUNs.
5242925	The limit for the number of concurrent delete jobs has been reached.
5242927	Unable to find qtree.
5242955	The UUID of the volume is required.
5242957	Failed to delete qtree with ID in the volume and SVM.

Error Code	Description
5242965	Invalid qtree path. The volume name component of the qtree path, must be the same as the volume specified with the parameter.
10485796	Cannot delete qtree because it contains a Storage Level Access Guard (SLAG).

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve qtree properties

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

Introduced In: 9.6

Retrieves properties for a specific qtree identified by the `volume.uuid` and the `id` in the API path.

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.

- `statistics.*`

Related ONTAP commands

- `qtree show`

Parameters

Name	Type	In	Required	Description
<code>volume.uuid</code>	string	path	True	Volume UUID
<code>id</code>	string	path	True	Qtree ID
<code>fields</code>	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
<code>_links</code>	_links	
<code>_tags</code>	array[string]	Tags are an optional way to track the uses of a resource. Tag values must be formatted as key:value strings.
<code>export_policy</code>	export_policy	Export Policy
<code>filesystem_path</code>	string	Path of the qtree directory. This path is relative to the volume root directory.
<code>group</code>	group	The user set as owner of the qtree.

Name	Type	Description
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
nas	nas	
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH. This field is to be deprecated and replaced with nas.path.
qos_policy	qos_policy	
security_style	string	Security style. Valid in POST or PATCH.
statistics	statistics	These are raw IOPS and throughput performance numbers. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
user	user	The user set as owner of the qtree.
volume	volume	Required in POST

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "filesystem_path": "/dir1/qtreen1",
  "group": {
    "id": "20001",
    "name": "unix_group1"
  },
  "id": 1,
  "nas": {
    "path": "/volume3/qtreen1"
  },
  "path": "/volume3/qtreen1",
  "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"
  },
  "security_style": "unix",
  "statistics": {
```

```
"iops_raw": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"status": "ok",
"throughput_raw": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 +0000"
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"unix_permissions": 493,
"user": {
  "id": "10001",
  "name": "unix_user1"
},
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
918235	A volume with UUID was not found.
2621462	The specified SVM does not exist.
5242889	Failed to get the qtree from volume.
5242956	Failed to obtain a qtree with ID.
5242965	Invalid qtree path. The volume name component of the qtree path, must be the same as the volume specified with the parameter.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

group

The user set as owner of the qtree.

Name	Type	Description
id	string	The numeric ID of the group that owns the qtree. Valid in POST or PATCH.
name	string	Alphanumeric group name of group that owns the qtree. Valid in POST or PATCH.

nas

Name	Type	Description
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.

qos_policy

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

iops_raw

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

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

throughput_raw

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

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

statistics

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

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

svm

Required in POST

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

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

user

The user set as owner of the qtree.

Name	Type	Description
id	string	The numeric ID of the user who owns the qtree. Valid in POST or PATCH.
name	string	Alphanumeric username of user who owns the qtree. Valid in POST or PATCH.

volume

Required in POST

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Update properties for a qtree

PATCH /storage/qtrees/{volume.uuid}/{id}

Introduced In: 9.6

Updates properties for a specific qtree.

Related ONTAP commands

- `qtree modify`
- `qtree rename`

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
id	string	path	True	Qtree ID

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

Request Body

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.
export_policy	export_policy	Export Policy
filesystem_path	string	Path of the qtree directory. This path is relative to the volume root directory.

Name	Type	Description
group	group	The user set as owner of the qtree.
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
nas	nas	
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH. This field is to be deprecated and replaced with nas.path.
qos_policy	qos_policy	
security_style	string	Security style. Valid in POST or PATCH.
statistics	statistics	These are raw IOPS and throughput performance numbers. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
user	user	The user set as owner of the qtree.
volume	volume	Required in POST

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "filesystem_path": "/dir1/qtreen1",
  "group": {
    "id": "20001",
    "name": "unix_group1"
  },
  "id": 1,
  "nas": {
    "path": "/volume3/qtreen1"
  },
  "path": "/volume3/qtreen1",
  "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"
  },
  "security_style": "unix",
  "statistics": {
```



```
"iops_raw": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"status": "ok",
"throughput_raw": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 +0000"
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"unix_permissions": 493,
"user": {
  "id": "10001",
  "name": "unix_user1"
},
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
262196	The request contains a field which cannot be set in this operation.
262278	Required field is missing in the request.
917505	Vserver not found.
917525	The specified volume does not exist in Vserver.
918235	A volume with UUID was not found.
5242887	Failed to modify qtree.
5242897	This operation is not permitted on read-only volume.
5242898	This operation is only permitted on a data Vserver.
5242902	Missing inputs.
5242915	Failed to assign qtree export policy to qtree.
5242927	Unable to find qtree.
5242945	Failed to modify qtree.
5242951	Export policy supplied does not belong to the specified export policy ID.

Error Code	Description
5242954	Failed to get the qtree from volume.
5242955	The UUID of the volume is required.
5242956	Failed to obtain a qtree with ID.
5242957	Failed to delete the qtree.
5242958	Failed to rename the qtree with ID in the volume and SVM.
5242959	Successfully renamed qtree but the modify operation failed.
5242965	Invalid qtree path. The volume name component of the qtree path, must be the same as the volume specified with the parameter.
5242967	UNIX user or group ID must be 32-bit unsigned integer.
5242971	Qtree was renamed. However, the path modification failed.
5242972	Cannot rename qtree as that name already exists on a volume in the Vserver.
5242973	Cannot rename qtree to name with path concurrently on volume in Vserver, unless non-root qtrees in enabled on the volume.
5242974	Moved qtree. However, other properties were not modified.
5242975	Renamed qtree and moved the qtree. However, other properties were not modified.
6622064	Security-style ntfs is not supported on SnapMirror Business Continuity (SMBC) relationship volume.
8454348	QoS on qtrees is not supported because not all nodes in the cluster can support it.
9437324	The security style unified is not supported.
23724050	Failed to resolve user or group name.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

group

The user set as owner of the qtree.

Name	Type	Description
id	string	The numeric ID of the group that owns the qtree. Valid in POST or PATCH.
name	string	Alphanumeric group name of group that owns the qtree. Valid in POST or PATCH.

nas

Name	Type	Description
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.

qos_policy

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

iops_raw

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

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

throughput_raw

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

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

statistics

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

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

svm

Required in POST

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

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

user

The user set as owner of the qtree.

Name	Type	Description
id	string	The numeric ID of the user who owns the qtree. Valid in POST or PATCH.
name	string	Alphanumeric username of user who owns the qtree. Valid in POST or PATCH.

volume

Required in POST

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

qtree

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

Name	Type	Description
_links	_links	

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.
export_policy	export_policy	Export Policy
filesystem_path	string	Path of the qtree directory. This path is relative to the volume root directory.
group	group	The user set as owner of the qtree.
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
nas	nas	
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH. This field is to be deprecated and replaced with nas.path.
qos_policy	qos_policy	
security_style	string	Security style. Valid in POST or PATCH.
statistics	statistics	These are raw IOPS and throughput performance numbers. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.

Name	Type	Description
user	user	The user set as owner of the qtree.
volume	volume	Required in POST

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Storage quota

Storage quota endpoint overview

Overview

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

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

Quota policy rule APIs

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

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

– POST /api/storage/quota/rules

– GET /api/storage/quota/rules

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

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

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

Enabling and disabling quotas

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

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

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

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

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

Quota report APIs

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

– GET /api/storage/quota/reports

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

Quota resize

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

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

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

– GET /api/storage/quota/reports

Manage storage quota reports

Storage quota reports endpoint overview

Overview

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

Quota report APIs

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

– GET /api/storage/quota/reports

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

Examples

Retrieving all the quota report records

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

The following example shows how to retrieve quota report records for all FlexVol volumes and FlexGroup volumes.

```
# The API:
GET /api/storage/quota/reports

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

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svml",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
```

```

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

```

```

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

```

```

    "href": "/api/storage/volumes/314a328f-502d-11e9-8771-005056a7f717"
  }
}
},
"index": 4611686018427387904,
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-005056a7f717/4611686018427387904"
  }
}
},
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "314a328f-502d-11e9-8771-005056a7f717",
    "name": "fg",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/314a328f-502d-11e9-8771-005056a7f717"
      }
    }
  },
  "index": 5764607523034234880,
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-005056a7f717/5764607523034234880"
    }
  }
}
},
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {

```



```

        "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
        }
    },
    "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
            "self": {
                "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
            }
        }
    },
    "index": 0,
    "_links": {
        "self": {
            "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/0"
        }
    }
},
{
    "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svml",
        "_links": {
            "self": {
                "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
            }
        }
    },
    "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
            "self": {
                "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
            }
        }
    },
    "index": 281474976710656,
    "_links": {
        "self": {

```

```
    "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-005056a7f717/281474976710656"
  }
}
},
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
      }
    }
  },
  "index": 1152921504606846976,
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-005056a7f717/1152921504606846976"
    }
  }
},
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
```

```

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

```

```

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

```

```

    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-005056a7f717/3459045988797251584"
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
          }
        }
      },
      "index": 4611686018427387904,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-005056a7f717/4611686018427387904"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {

```

```
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
      }
    }
  },
  "index": 4611967493404098560,
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-005056a7f717/4611967493404098560"
    }
  }
},
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
      }
    }
  },
  "index": 5764607523034234880,
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-005056a7f717/5764607523034234880"
    }
  }
}
],
```

```
"num_records": 15,
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/"
  }
}
}
```

Retrieving a specific quota report record

This API is used to retrieve a specific quota report record.

The following example shows how to retrieve a single quota report user record.

```
# The API:
GET /api/storage/quota/reports/{volume.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/281474976710656" -H 'accept: application/hal+json'

# Response for quota report user record:
{
  "svm": {
    "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
      }
    }
  },
  "volume": {
    "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
      }
    }
  },
  "index": 281474976710656,
  "type": "user",
```

```
"users": [
  {
    "name": "fred",
    "id" : "300008"
  }
],
"qtree": {
  "name": "qt1",
  "id": 1,
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
    }
  }
},
"space": {
  "hard_limit": 41943040,
  "soft_limit": 31457280,
  "used": {
    "total": 10567680,
    "soft_limit_percent": 34,
    "hard_limit_percent": 25
  }
},
"files": {
  "soft_limit": 30,
  "hard_limit": 40,
  "used": {
    "total": 11,
    "soft_limit_percent": 37,
    "hard_limit_percent": 28
  }
},
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/281474976710656"
  }
}
}
```

Retrieving a single quota report multi-user record

```

# The API:
GET /api/storage/quota/reports/{volume.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/281474976710656" -H 'accept: application/hal+json'

# Response for quota report multi-user record:
{
  "svm": {
    "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
      }
    }
  },
  "volume": {
    "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
      }
    }
  },
  "index": 1153484454560268288,
  "type": "user",
  "users": [
    {
      "name": "fred",
      "id" : "300008"
    },
    {
      "name": "john",
      "id" : "300009"
    },
    {
      "name": "smith",
      "id" : "300010"
    }
  ],
  "qtree": {
    "name": "qt1",
    "id": 1,

```

```

    "_links": {
      "self": {
        "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
      }
    }
  },
  "space": {
    "hard_limit": 41943040,
    "soft_limit": 31457280,
    "used": {
      "total": 10567680,
      "soft_limit_percent": 34,
      "hard_limit_percent": 25
    }
  },
  "files": {
    "soft_limit": 30,
    "hard_limit": 40,
    "used": {
      "total": 11,
      "soft_limit_percent": 37,
      "hard_limit_percent": 28
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/1153484454560268288"
    }
  }
}

```

Retrieving a single quota report group record

```

# The API:
GET /api/storage/quota/reports/{volume.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/3459045988797251584" -H 'accept: application/hal+json'

```

```
# Response for quota report group record:
{
  "svm": {
    "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
      }
    }
  },
  "volume": {
    "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
      }
    }
  },
  "index": 3459045988797251584,
  "type": "group",
  "group": {
    "name" : "test_group",
    "id"   : "500009"
  },
  "qtree": {
    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
      }
    }
  },
  "space": {
    "hard_limit": 41943040,
    "soft_limit": 31457280,
    "used": {
      "total": 10567680,
      "soft_limit_percent": 34,
      "hard_limit_percent": 25
    }
  },
  "files": {
    "soft_limit": 30,

```

```
"hard_limit": 40,
"used": {
  "total": 11,
  "soft_limit_percent": 37,
  "hard_limit_percent": 28
}
},
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/3459045988797251584"
  }
}
}
```

Retrieving a single quota report tree record

```
# The API:
GET /api/storage/quota/reports/{volume.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/4612248968380809216" -H 'accept: application/hal+json'

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

```
    }
  },
  "index": 4612248968380809216,
  "type": "tree",
  "qtree": {
    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
      }
    }
  },
  "space": {
    "hard_limit": 41943040,
    "soft_limit": 31457280,
    "used": {
      "total": 10567680,
      "soft_limit_percent": 34,
      "hard_limit_percent": 25
    }
  },
  "files": {
    "soft_limit": 30,
    "hard_limit": 40,
    "used": {
      "total": 11,
      "soft_limit_percent": 37,
      "hard_limit_percent": 28
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/4612248968380809216"
    }
  }
}
```

Retrieving only records enforced by non-default rules

```

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/quota/reports?show_default_records=false" -H 'accept:
application/hal+json'

# Response from only non-default records
{
"records": [
  {
    "svm": {
      "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
        }
      }
    },
    "volume": {
      "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-
005056a7657c"
        }
      }
    },
    "index": 4612248968380809216,
    "type": "tree",
    "qtree": {
      "name": "qt1",
      "id": 1,
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-
005056a7657c/1"
        }
      }
    },
    "space": {
      "hard_limit": 41943040,
      "soft_limit": 31457280,
      "used": {
        "total": 10567680,
        "soft_limit_percent": 34,

```

```

    "hard_limit_percent": 25
  }
},
"files": {
  "soft_limit": 30,
  "hard_limit": 40,
  "used": {
    "total": 11,
    "soft_limit_percent": 37,
    "hard_limit_percent": 28
  }
},
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/4612248968380809216"
  }
}
},
{
  "svm": {
    "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
      }
    }
  },
  "volume": {
    "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
      }
    }
  },
  "index": 1153484454560268288,
  "type": "user",
  "users": [
    {
      "name": "fred",
      "id": "300008"
    }
  ],

```

```

    {
      "name": "john",
      "id" : "300009"
    },
    {
      "name": "smith",
      "id" : "300010"
    }
  ],
  "qtree": {
    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
      }
    }
  },
  "space": {
    "hard_limit": 41943040,
    "soft_limit": 31457280,
    "used": {
      "total": 10567680,
      "soft_limit_percent": 34,
      "hard_limit_percent": 25
    }
  },
  "files": {
    "soft_limit": 30,
    "hard_limit": 40,
    "used": {
      "total": 11,
      "soft_limit_percent": 37,
      "hard_limit_percent": 28
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/1153484454560268288"
    }
  }
},
"num_records": 2,

```



```
"_links": {
  "self": {
    "href": "/api/storage/quota/reports?show_default_records=false"
  }
}
}
```

Retrieving quota report records with query parameters

The following example shows how to retrieve tree type quota report records.

```
# The API:
GET /api/storage/quota/reports

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports?type=tree" -H
'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "903e54ee-6ccf-11ea-bc35-005056823577",
        "name": "svml",
        "_links": {
          "self": {
            "href": "/api/svm/svms/903e54ee-6ccf-11ea-bc35-005056823577"
          }
        }
      },
      "volume": {
        "uuid": "8812b000-6e1e-11ea-9bad-00505682cd5c",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/8812b000-6e1e-11ea-9bad-00505682cd5c"
          }
        }
      },
      "index": 2305843013508661248,
      "type": "tree",

```

```

    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/8812b000-6e1e-11ea-9bad-
00505682cd5c/2305843013508661248"
      }
    }
  },
  {
    "svm": {
      "uuid": "903e54ee-6ccf-11ea-bc35-005056823577",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/903e54ee-6ccf-11ea-bc35-005056823577"
        }
      }
    },
    "volume": {
      "uuid": "a5ceebd2-6ccf-11ea-bc35-005056823577",
      "name": "fg",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/a5ceebd2-6ccf-11ea-bc35-
005056823577"
        }
      }
    },
    "index": 2305843013508661248,
    "type": "tree",
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/a5ceebd2-6ccf-11ea-bc35-
005056823577/2305843013508661248"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/storage/quota/reports?type=tree"
  }
}
}

```

Retrieving all the quota reports of a specific volume and the files fields

```
# The API:
GET /api/storage/quota/reports

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/quota/reports?volume.name=fv&fields=files" -H 'accept:
application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "903e54ee-6ccf-11ea-bc35-005056823577",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/903e54ee-6ccf-11ea-bc35-005056823577"
          }
        }
      },
      "volume": {
        "uuid": "8812b000-6e1e-11ea-9bad-00505682cd5c",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/8812b000-6e1e-11ea-9bad-
00505682cd5c"
          }
        }
      },
      "index": 410328290557952,
      "files": {
        "soft_limit": 20,
        "hard_limit": 30,
        "used": {
          "total": 0,
          "soft_limit_percent": 0,
          "hard_limit_percent": 0
        }
      },
      "_links": {
        "self": {

```

```

    "href": "/api/storage/quota/reports/8812b000-6e1e-11ea-9bad-
00505682cd5c/410328290557952"
  }
}
},
{
  "svm": {
    "uuid": "903e54ee-6ccf-11ea-bc35-005056823577",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/903e54ee-6ccf-11ea-bc35-005056823577"
      }
    }
  },
  "volume": {
    "uuid": "8812b000-6e1e-11ea-9bad-00505682cd5c",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/8812b000-6e1e-11ea-9bad-
00505682cd5c"
      }
    }
  },
  "index": 2305843013508661248,
  "files": {
    "soft_limit": 200,
    "hard_limit": 400,
    "used": {
      "total": 4,
      "soft_limit_percent": 2,
      "hard_limit_percent": 1
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/8812b000-6e1e-11ea-9bad-
00505682cd5c/2305843013508661248"
    }
  }
},
"num_records": 2,
"_links": {
  "self": {

```

```

    "href": "/api/storage/quota/reports?volume.name=fv&fields=files"
  }
}
}

```

Retrieve the quota report records for all FlexVol and FlexGroup volumes

GET /storage/quota/reports

Introduced In: 9.6

Retrieves the quota report records for all FlexVol volumes and FlexGroup volumes.

Related ONTAP commands

- `quota report`

Parameters

Name	Type	In	Required	Description
files.soft_limit	integer	query	False	Filter by files.soft_limit
files.used.total	integer	query	False	Filter by files.used.total
files.used.soft_limit_percent	integer	query	False	Filter by files.used.soft_limit_percent
files.used.hard_limit_percent	integer	query	False	Filter by files.used.hard_limit_percent
files.hard_limit	integer	query	False	Filter by files.hard_limit
users.id	string	query	False	Filter by users.id
users.name	string	query	False	Filter by users.name
volume.uuid	string	query	False	Filter by volume.uuid
volume.name	string	query	False	Filter by volume.name

Name	Type	In	Required	Description
type	string	query	False	Filter by type
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
space.used.soft_limit_percent	integer	query	False	Filter by space.used.soft_limit_percent
space.used.total	integer	query	False	Filter by space.used.total
space.used.hard_limit_percent	integer	query	False	Filter by space.used.hard_limit_percent
space.hard_limit	integer	query	False	Filter by space.hard_limit
space.soft_limit	integer	query	False	Filter by space.soft_limit
specifier	string	query	False	Filter by specifier
qtree.name	string	query	False	Filter by qtree.name
qtree.id	integer	query	False	Filter by qtree.id
group.name	string	query	False	Filter by group.name
group.id	string	query	False	Filter by group.id
index	integer	query	False	Filter by index
show_default_records	boolean	query	False	<p>The default is true for GET calls. When set to false, the default records are not reported.</p> <ul style="list-style-type: none"> • Introduced in: 9.7 • Default value: 1

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

Response

Status: 200, Ok

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "files": {
      "hard_limit": 0,
      "soft_limit": 0,
      "used": {
        "hard_limit_percent": 0,
        "soft_limit_percent": 0,
        "total": 0
      }
    },
    "group": {
      "id": "string",
      "name": "string"
    },
    "index": 0,
    "qtree": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 1,
      "name": "qt1"
    },
    "space": {
      "hard_limit": 0,
      "soft_limit": 0,
      "used": {
        "hard_limit_percent": 0,
```



```

    "soft_limit_percent": 0,
    "total": 0
  }
},
"specifier": "string",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"type": "tree",
"users": {
  "id": "string",
  "name": "string"
},
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

used

Name	Type	Description
hard_limit_percent	integer	Total files used as a percentage of file hard limit
soft_limit_percent	integer	Total files used as a percentage of file soft limit
total	integer	Total files used

files

Name	Type	Description
hard_limit	integer	File hard limit
soft_limit	integer	File soft limit
used	used	

group

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

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

qtree

This parameter specifies the target qtree associated with the user, group, or tree record. For a user/group quota policy rule at volume level, this parameter is not valid. For a default tree quota policy rule, this parameter is specified as "" or "*". For a tree quota policy rule at qtree level, this parameter specifies a qtree name and a qtree identifier.

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

used

Name	Type	Description
hard_limit_percent	integer	Total space used as a percentage of space hard limit
soft_limit_percent	integer	Total space used as a percentage of space soft limit
total	integer	Total space used

space

Name	Type	Description
hard_limit	integer	Space hard limit in bytes
soft_limit	integer	Space soft limit in bytes
used	used	

svm

Name	Type	Description
_links	_links	

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

users

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

volume

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

quota_report

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

Name	Type	Description
index	integer	Index that identifies a unique quota record. Valid in URL.
qtree	qtree	This parameter specifies the target qtree associated with the user, group, or tree record. For a user/group quota policy rule at volume level, this parameter is not valid. For a default tree quota policy rule, this parameter is specified as "" or "*". For a tree quota policy rule at qtree level, this parameter specifies a qtree name and a qtree identifier.
space	space	
specifier	string	Quota specifier
svm	svm	
type	string	Quota type associated with the quota record.
users	array[users]	This parameter specifies the target user or users associated with the given quota report record. This parameter is available for user quota records and is not available for group or tree quota records. The target user or users are identified by a user name and user identifier. The user name can be a UNIX user name or a Windows user name, and the identifier can be a UNIX user identifier or a Windows security identifier.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a specific quota report record

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

Introduced In: 9.6

Retrieves a specific quota report record.

Related ONTAP commands

- `quota report`

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
index	integer	path	True	Quota report index
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
files	files	

Name	Type	Description
group	group	This parameter specifies the target group associated with the given quota report record. This parameter is available for group quota records and is not available for user or tree quota records. The target group is identified by a UNIX group name and UNIX group identifier.
index	integer	Index that identifies a unique quota record. Valid in URL.
qtree	qtree	This parameter specifies the target qtree associated with the user, group, or tree record. For a user/group quota policy rule at volume level, this parameter is not valid. For a default tree quota policy rule, this parameter is specified as "" or "*". For a tree quota policy rule at qtree level, this parameter specifies a qtree name and a qtree identifier.
space	space	
specifier	string	Quota specifier
svm	svm	
type	string	Quota type associated with the quota record.
users	array[users]	This parameter specifies the target user or users associated with the given quota report record. This parameter is available for user quota records and is not available for group or tree quota records. The target user or users are identified by a user name and user identifier. The user name can be a UNIX user name or a Windows user name, and the identifier can be a UNIX user identifier or a Windows security identifier.
volume	volume	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "files": {
    "hard_limit": 0,
    "soft_limit": 0,
    "used": {
      "hard_limit_percent": 0,
      "soft_limit_percent": 0,
      "total": 0
    }
  },
  "group": {
    "id": "string",
    "name": "string"
  },
  "index": 0,
  "qtree": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 1,
    "name": "qt1"
  },
  "space": {
    "hard_limit": 0,
    "soft_limit": 0,
    "used": {
      "hard_limit_percent": 0,
      "soft_limit_percent": 0,
      "total": 0
    }
  },
  "specifier": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  }
}
```

```

    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "tree",
  "users": {
    "id": "string",
    "name": "string"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}

```

Error

Status: Default

ONTAP Error Response Codes

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

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

used

Name	Type	Description
hard_limit_percent	integer	Total files used as a percentage of file hard limit
soft_limit_percent	integer	Total files used as a percentage of file soft limit
total	integer	Total files used

files

Name	Type	Description
hard_limit	integer	File hard limit
soft_limit	integer	File soft limit
used	used	

group

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

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

qtree

This parameter specifies the target qtree associated with the user, group, or tree record. For a user/group

quota policy rule at volume level, this parameter is not valid. For a default tree quota policy rule, this parameter is specified as "" or "*". For a tree quota policy rule at qtree level, this parameter specifies a qtree name and a qtree identifier.

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

used

Name	Type	Description
hard_limit_percent	integer	Total space used as a percentage of space hard limit
soft_limit_percent	integer	Total space used as a percentage of space soft limit
total	integer	Total space used

space

Name	Type	Description
hard_limit	integer	Space hard limit in bytes
soft_limit	integer	Space soft limit in bytes
used	used	

svm

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

users

Name	Type	Description
id	string	Quota target user ID

Name	Type	Description
name	string	Quota target user name

volume

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage storage quota policy rules

Storage quota rules endpoint overview

Overview

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

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

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

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

Quota policy rule APIs

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

– POST /api/storage/quota/rules

– GET /api/storage/quota/rules

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

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

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

Examples

Retrieving all quota policy rules

This API is used to retrieve all quota policy rules.

The following example shows how to retrieve quota policy rules for FlexVol volumes and FlexGroup volumes.

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

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

# The response:
{
  "records": [
    {
      "svm": {
```

```

    "uuid": "038545f8-9ff8-11e8-bce6-005056a73bed",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/038545f8-9ff8-11e8-bce6-005056a73bed"
      }
    }
  },
  "volume": {
    "uuid": "ab3df793-0f02-43c6-9514-4f142fc8cc92",
    "name": "voll1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/ab3df793-0f02-43c6-9514-4f142fc8cc92"
      }
    }
  },
  "uuid": "66319cbe-b837-11e8-9c5a-005056a7e88c",
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules/66319cbe-b837-11e8-9c5a-005056a7e88c"
    }
  }
},
{
  "svm": {
    "uuid": "038545f8-9ff8-11e8-bce6-005056a73bed",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/038545f8-9ff8-11e8-bce6-005056a73bed"
      }
    }
  },
  "volume": {
    "uuid": "ab3df793-0f02-43c6-9514-4f142fc8cc92",
    "name": "voll1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/ab3df793-0f02-43c6-9514-4f142fc8cc92"
      }
    }
  }
},

```



```

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

```

Retrieving a specific quota policy rule

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

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

```

# The API:
GET /api/storage/quota/rules/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/264a9e0b-2e03-11e9-
a610-005056a7b72d' -H 'accept: application/hal+json'

# Response for a user rule at a qtree level:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",

```

```

    "name": "voll1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/clb64eea-ca8b-45ec-9397-
ab489830d268"
      }
    }
  },
  "uuid": "264a9e0b-2e03-11e9-a610-005056a7b72d",
  "type": "user",
  "users": [ {"name" : "fred"} ],
  "qtree": {
    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/clb64eea-ca8b-45ec-9397-
ab489830d268/1"
      }
    }
  },
  "user_mapping": on,
  "space": {
    "hard_limit": 1222800,
    "soft_limit": 51200
  },
  "files": {
    "hard_limit": 100,
    "soft_limit": 80
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules/264a9e0b-2e03-11e9-a610-
005056a7b72d"
    }
  }
}

```

Retrieving a quota policy multi-user rule at the volume level

```

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/0ab84fba-19aa-11e9-

```

```
a04d-005056a72f42' -H 'accept: application/hal+json'

# Response for a multi-user rule at volume level:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "voll",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-
ab489830d268"
      }
    }
  },
  "uuid": "0ab84fba-19aa-11e9-a04d-005056a72f42",
  "type": "user",
  "users": [
    {
      "name": "sam",
    },
    {
      "name": "smith",
    },
    {
      "id": "300010",
    },
  ],
  "space": {
    "hard_limit": 1222800,
    "soft_limit": 51200
  },
  "files": {
    "hard_limit": 100,
    "soft_limit": 80
  },
  "_links": {
    "self": {
```

```
"href": "/api/storage/quota/rules/0ab84fba-19aa-11e9-a04d-005056a72f42"
  }
}
```

Retrieving a quota policy default tree rule

```
# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/4a276b8c-1753-11e9-8101-005056a760e0' -H 'accept: application/hal+json'

# Response for a default tree rule:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "voll1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-ab489830d268"
      }
    }
  },
  "uuid": "4a276b8c-1753-11e9-8101-005056a760e0",
  "type": "tree",
  "qtree": {
    "name": ""
  },
  "space": {
    "hard_limit": 1034000,
    "soft_limit": 51200
  },
  "files": {
    "hard_limit": 20,
    "soft_limit": 10
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules/4a276b8c-1753-11e9-8101-005056a760e0"
    }
  }
}
```

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

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

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

Retrieving a quota policy group rule for a specific qtree

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

# Response for a group rule:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "voll1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-ab489830d268"
      }
    }
  },
  "uuid": "b9236852-19ab-11e9-a04d-005056a72f42",
  "type": "group",
  "group": {"name" : "group1"},
  "qtree": {
    "name": "qt1",
```

```

    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/clb64eea-ca8b-45ec-9397-
ab489830d268/1"
      }
    },
    "space": {
      "hard_limit": 2097152,
      "soft_limit": 1572864
    },
    "files": {
      "hard_limit": 250,
      "soft_limit": 200
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/rules/b9236852-19ab-11e9-a04d-
005056a72f42"
      }
    }
  }
}

```

Creating a quota policy rule

This API is used to create a new quota policy rule. When an explicit rule or a qtree-scoped rule of a type is created on a volume, a default rule of the same type is automatically added if it does not already exist on the volume.

The following example shows how to create a quota policy user rule using POST.

```

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

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

```



```

},
"volume": {
  "name": "voll1"
},
"type": "user",
"users": [ {"name" : "jsmith"} ],
"qtree": {
  "name": "qt1"
},
"user_mapping": "on",
"space": {
  "hard_limit": 8192,
  "soft_limit": 1024
},
"files": {
  "hard_limit": 20,
  "soft_limit": 10
}
}

# The response
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svml1"
      },
      "volume": {
        "name": "fv"
      },
      "uuid": "3220eea6-5049-11e9-bfb7-005056a7f717",
      "type": "user",
      "users": [
        {
          "name" : "jsmith"
        }
      ],
      "qtree": {
        "name": "qt1"
      },
      "user_mapping": "on",
      "space": {
        "hard_limit": 8192,
        "soft_limit": 1024
      },
    }
  ]
}

```

```

    "files": {
      "hard_limit": 20,
      "soft_limit": 10
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/rules/3220eea6-5049-11e9-bfb7-005056a7f717"
      }
    }
  ],
  "job": {
    "uuid": "32223924-5049-11e9-bfb7-005056a7f717",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/32223924-5049-11e9-bfb7-005056a7f717"
      }
    }
  }
}

```

Creating a quota policy group rule using POST.

```

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

# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/quota/rules?return_records=true' -H 'accept: application/hal+json' -d @test_quota_post.txt
test_quota_post.txt (body):
{
  "svm": {
    "name": "svm1"
  },
  "volume": {
    "name": "voll1"
  },
  "type": "group",
  "group": {
    "name": "test_group1"
  }
}

```

```

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

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

```

```

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

```

Creating a quota policy tree rule using POST

```

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

# The call:
curl -X POST 'https://<mgmt-
ip>/api/storage/quota/rules?return_records=true' -H 'accept:
application/hal+json' -d @test_quota_post.txt
test_quota_post.txt (body):
{
  "svm": {
    "name": "svm1"
  },
  "volume": {
    "name": "voll1"
  },
  "type": "tree",
  "qtree": {
    "name": "qt1"
  },
  "space": {
    "hard_limit": 8192,
    "soft_limit": 1024
  },
  "files": {
    "hard_limit": 20,
    "soft_limit": 10
  }
}

```

```

}

# The response
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1"
      },
      "volume": {
        "name": "fv"
      },
      "uuid": "e5eb03be-504a-11e9-bfb7-005056a7f717",
      "type": "tree",
      "qtree": {
        "name": "qt1"
      },
      "space": {
        "hard_limit": 8192,
        "soft_limit": 1024
      },
      "files": {
        "hard_limit": 20,
        "soft_limit": 10
      },
      "_links": {
        "self": {
          "href": "/api/storage/quota/rules/e5eb03be-504a-11e9-bfb7-
005056a7f717"
        }
      }
    },
    {
      "job": {
        "uuid": "32223924-5049-11e9-bfb7-005056a7f717",
        "_links": {
          "self": {
            "href": "/api/cluster/jobs/32223924-5049-11e9-bfb7-005056a7f717"
          }
        }
      }
    }
  ]
}

```

Updating the quota policy rule

This API is used to update a quota policy rule.

The following example shows how to update a quota policy rule.

```
# The API:
PATCH /storage/quota/rules/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/quota/rules/364d38eb-8e87-11e8-a806-005056a7e73a" -H 'accept: application/hal+json' -d
"@test_quota_patch.txt"
test_quota_patch.txt(body):
{
  "space": {
    "hard_limit": 16554,
    "soft_limit": 8192
  },
  "files": {
    "hard_limit": 40,
    "soft_limit": 20
  }
}
```

Deleting the quota policy rule

This API is used to delete a quota policy rule.

The following example shows how to delete a quota policy rule.

```
# The API:
DELETE /storage/quota/rules/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/quota/rules/364d38eb-8e87-11e8-a806-005056a7e73a" -H 'accept: application/hal+json'
```

Retrieve quota policy rules for all FlexVol and FlexGroup volumes

GET /storage/quota/rules

Introduced In: 9.6

Retrieves quota policy rules configured for all FlexVol volumes and FlexGroup volumes.

Related ONTAP commands

- quota policy rule show

Parameters

Name	Type	In	Required	Description
type	string	query	False	Filter by type
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
space.hard_limit	integer	query	False	Filter by space.hard_limit
space.soft_limit	integer	query	False	Filter by space.soft_limit
qtree.name	string	query	False	Filter by qtree.name
qtree.id	integer	query	False	Filter by qtree.id
group.id	string	query	False	Filter by group.id
group.name	string	query	False	Filter by group.name
files.soft_limit	integer	query	False	Filter by files.soft_limit
files.hard_limit	integer	query	False	Filter by files.hard_limit
users.id	string	query	False	Filter by users.id
users.name	string	query	False	Filter by users.name
volume.uuid	string	query	False	Filter by volume.uuid

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

Response

Status: 200, Ok

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "qtree": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 1,
      "name": "qt1"
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "type": "tree",
    "users": {
    },
    "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  },
}
```

```
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

files

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

group

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

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

qtree

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

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

space

Name	Type	Description
hard_limit	integer	This parameter specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.
soft_limit	integer	This parameter specifies the space soft limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.

svm

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

users

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

volume

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

quota_rule

Name	Type	Description
_links	_links	
files	files	

Name	Type	Description
group	group	<p>This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.</p>
qtree	qtree	<p>This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.</p>
space	space	
svm	svm	
type	string	<p>This parameter specifies the quota policy rule type. This is required in POST only and can take either one of the "user", "group" or "tree" values.</p>

Name	Type	Description
user_mapping	boolean	This parameter enables user mapping for user quota policy rules. This is valid in POST or PATCH for user quota policy rules only.
users	array[users]	This parameter specifies the target user to which the user quota policy rule applies. This parameter takes single or multiple user names or identifiers. This parameter is valid only for the POST operation of a user quota policy rule. If this parameter is used as an input to create a group or a tree quota policy rule, the POST operation will fail with an appropriate error. For POST, this input parameter takes either a user name or a user identifier, not both. For default quota rules, the user name must be chosen and specified as "". For explicit user quota rules, this parameter can indicate either a user name or user identifier. The user name can be a UNIX user name or a Windows user name. If a name contains a space, enclose the entire value in quotes. A UNIX user name cannot include a backslash () or an @ sign; user names with these characters are treated as Windows names. The user identifier can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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

POST /storage/quota/rules

Introduced In: 9.6

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

Important notes:

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

Required properties

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

Recommended optional properties

- `space.hard_limit` - Specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes.

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

Related ONTAP commands

- `quota policy rule create`

Parameters

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

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

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

Name	Type	Description
qtree	qtree	This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.
space	space	
svm	svm	
type	string	This parameter specifies the quota policy rule type. This is required in POST only and can take either one of the "user", "group" or "tree" values.
user_mapping	boolean	This parameter enables user mapping for user quota policy rules. This is valid in POST or PATCH for user quota policy rules only.

Name	Type	Description
users	array[users]	<p>This parameter specifies the target user to which the user quota policy rule applies. This parameter takes single or multiple user names or identifiers. This parameter is valid only for the POST operation of a user quota policy rule. If this parameter is used as an input to create a group or a tree quota policy rule, the POST operation will fail with an appropriate error. For POST, this input parameter takes either a user name or a user identifier, not both. For default quota rules, the user name must be chosen and specified as "". For explicit user quota rules, this parameter can indicate either a user name or user identifier. The user name can be a UNIX user name or a Windows user name. If a name contains a space, enclose the entire value in quotes. A UNIX user name cannot include a backslash () or an @ sign; user names with these characters are treated as Windows names. The user identifier can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.</p>
uuid	string	<p>Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.</p>
volume	volume	

Example request

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

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Headers

Name	Description	Type
Location	Useful for tracking the resource location	string

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
917927	The specified volume was not found.
918232	Either <code>volume.name</code> or <code>volume.uuid</code> must be provided.
918236	The specified <code>volume.uuid</code> and <code>volume.name</code> refer to different volumes.
2621462	The specified SVM does not exist.
2621706	The specified <code>svm.uuid</code> and <code>svm.name</code> do not refer to the same SVM.

Error Code	Description
2621707	No SVM was specified. Either <code>svm.name</code> or <code>svm.uuid</code> must be supplied.
5308501	Mapping from Windows user to UNIX user for user rule was unsuccessful.
5308502	Mapping from UNIX user to Windows user for user rule was unsuccessful.
5308552	Failed to get default quota policy name for SVM.
5308561	Failed to obtain volume quota state or invalid quota state obtained for volume.
5308562	<code>users</code> is a required input for creating a user rule and <code>group</code> is not allowed.
5308563	<code>group</code> is a required input for creating a group rule and <code>users</code> is not allowed.
5308564	<code>qtree.name</code> is a required input for creating a tree rule and <code>users</code> and <code>group</code> are not allowed.
5308565	Only one of <code>name</code> or <code>id</code> is allowed for each entry in the <code>users</code> array.
5308566	Only one of <code>name</code> or <code>id</code> is allowed for <code>group</code> .
5308568	Quota policy rule create operation succeeded, but quota resize failed due to internal error. To activate the rule, disable and enable quotas for this volume.
5308571	Quota policy rule create operation succeeded, but quota resize is skipped. To activate the rule, disable and enable quotas for this volume.
5308573	Input value is greater than limit for field.
5308574	Input value is out of range for field.
5308575	Input value is incorrectly larger than listed field.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

files

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

group

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

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

qtree

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

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

space

Name	Type	Description
hard_limit	integer	This parameter specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.
soft_limit	integer	This parameter specifies the space soft limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.

svm

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

users

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

volume

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

quota_rule

Name	Type	Description
_links	_links	
files	files	

Name	Type	Description
group	group	<p>This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.</p>
qtree	qtree	<p>This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.</p>
space	space	
svm	svm	
type	string	<p>This parameter specifies the quota policy rule type. This is required in POST only and can take either one of the "user", "group" or "tree" values.</p>

Name	Type	Description
user_mapping	boolean	This parameter enables user mapping for user quota policy rules. This is valid in POST or PATCH for user quota policy rules only.
users	array[users]	This parameter specifies the target user to which the user quota policy rule applies. This parameter takes single or multiple user names or identifiers. This parameter is valid only for the POST operation of a user quota policy rule. If this parameter is used as an input to create a group or a tree quota policy rule, the POST operation will fail with an appropriate error. For POST, this input parameter takes either a user name or a user identifier, not both. For default quota rules, the user name must be chosen and specified as "". For explicit user quota rules, this parameter can indicate either a user name or user identifier. The user name can be a UNIX user name or a Windows user name. If a name contains a space, enclose the entire value in quotes. A UNIX user name cannot include a backslash () or an @ sign; user names with these characters are treated as Windows names. The user identifier can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.
volume	volume	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Delete a quota policy rule

DELETE /storage/quota/rules/{uuid}

Introduced In: 9.7

Deletes a quota policy rule.

Related ONTAP commands

- `quota policy rule delete`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Rule UUID

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

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
5308545	The specified quota rule UUID is invalid.
5308561	Failed to obtain volume quota state or invalid quota state obtained for volume.
5308569	Quota policy rule delete operation succeeded, but quota resize failed due to internal error.
5308572	Quota policy rule delete operation succeeded, however the rule is still being enforced. To stop enforcing the rule, disable quotas and enable them again for this volume.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve quota policy rule properties

GET /storage/quota/rules/{uuid}

Introduced In: 9.7

Retrieves properties for a specific quota policy rule.

Related ONTAP commands

- `quota policy rule show`

Parameters

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

Response

Status: 200, Ok

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

Name	Type	Description
qtree	qtree	This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.
space	space	
svm	svm	
type	string	This parameter specifies the quota policy rule type. This is required in POST only and can take either one of the "user", "group" or "tree" values.
user_mapping	boolean	This parameter enables user mapping for user quota policy rules. This is valid in POST or PATCH for user quota policy rules only.

Name	Type	Description
users	array[users]	<p>This parameter specifies the target user to which the user quota policy rule applies. This parameter takes single or multiple user names or identifiers. This parameter is valid only for the POST operation of a user quota policy rule. If this parameter is used as an input to create a group or a tree quota policy rule, the POST operation will fail with an appropriate error. For POST, this input parameter takes either a user name or a user identifier, not both. For default quota rules, the user name must be chosen and specified as "". For explicit user quota rules, this parameter can indicate either a user name or user identifier. The user name can be a UNIX user name or a Windows user name. If a name contains a space, enclose the entire value in quotes. A UNIX user name cannot include a backslash () or an @ sign; user names with these characters are treated as Windows names. The user identifier can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.</p>
uuid	string	<p>Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.</p>
volume	volume	

Example response

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

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
5308544	The specified quota rule UUID is invalid.
5308545	Unable to retrieve rule for the specified quota rule UUID.
5308576	Parameter <code>show_default_records</code> only allowed for GET collection.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

files

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

group

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

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

qtree

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

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

space

Name	Type	Description
hard_limit	integer	This parameter specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.
soft_limit	integer	This parameter specifies the space soft limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.

svm

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

users

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

volume

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Update quota policy rule properties

PATCH /storage/quota/rules/{uuid}

Introduced In: 9.7

Updates properties of a specific quota policy rule.

Important notes:

- The quota resize functionality is supported with the PATCH operation.
- Quota resize allows you to modify the quota limits, directly in the filesystem.
- The quota must be enabled on a FlexVol or a FlexGroup volume for `quota resize` to take effect.
- If the quota is disabled on the volume, the quota policy rule PATCH API modifies the rule, but this does not affect the limits in the filesystem.

Related ONTAP commands

- `quota policy rule modify`
- `quota resize`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Rule UUID

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

Request Body

Name	Type	Description
_links	_links	
files	files	

Name	Type	Description
group	group	<p>This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.</p>
qtree	qtree	<p>This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.</p>
space	space	
svm	svm	
type	string	<p>This parameter specifies the quota policy rule type. This is required in POST only and can take either one of the "user", "group" or "tree" values.</p>

Name	Type	Description
user_mapping	boolean	This parameter enables user mapping for user quota policy rules. This is valid in POST or PATCH for user quota policy rules only.
users	array[users]	This parameter specifies the target user to which the user quota policy rule applies. This parameter takes single or multiple user names or identifiers. This parameter is valid only for the POST operation of a user quota policy rule. If this parameter is used as an input to create a group or a tree quota policy rule, the POST operation will fail with an appropriate error. For POST, this input parameter takes either a user name or a user identifier, not both. For default quota rules, the user name must be chosen and specified as "". For explicit user quota rules, this parameter can indicate either a user name or user identifier. The user name can be a UNIX user name or a Windows user name. If a name contains a space, enclose the entire value in quotes. A UNIX user name cannot include a backslash () or an @ sign; user names with these characters are treated as Windows names. The user identifier can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.
volume	volume	

Example request

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

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
5308501	Mapping from Windows user to UNIX user for user rule was unsuccessful.
5308502	Mapping from UNIX user to Windows user for user rule was unsuccessful.
5308545	The specified quota rule UUID is invalid.
5308561	Failed to obtain volume quota state or invalid quota state obtained for volume.
5308567	Quota policy rule modify operation succeeded, but quota resize failed due to internal error.
5308573	Input value is greater than limit for field.
5308574	Input value is out of range for field.
5308575	Input value is incorrectly larger than listed field.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

files

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

group

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

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

qtree

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

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

space

Name	Type	Description
hard_limit	integer	This parameter specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.
soft_limit	integer	This parameter specifies the space soft limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.

svm

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

users

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

volume

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

quota_rule

Name	Type	Description
_links	_links	
files	files	

Name	Type	Description
group	group	<p>This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.</p>
qtree	qtree	<p>This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.</p>
space	space	
svm	svm	
type	string	<p>This parameter specifies the quota policy rule type. This is required in POST only and can take either one of the "user", "group" or "tree" values.</p>

Name	Type	Description
user_mapping	boolean	This parameter enables user mapping for user quota policy rules. This is valid in POST or PATCH for user quota policy rules only.
users	array[users]	This parameter specifies the target user to which the user quota policy rule applies. This parameter takes single or multiple user names or identifiers. This parameter is valid only for the POST operation of a user quota policy rule. If this parameter is used as an input to create a group or a tree quota policy rule, the POST operation will fail with an appropriate error. For POST, this input parameter takes either a user name or a user identifier, not both. For default quota rules, the user name must be chosen and specified as "". For explicit user quota rules, this parameter can indicate either a user name or user identifier. The user name can be a UNIX user name or a Windows user name. If a name contains a space, enclose the entire value in quotes. A UNIX user name cannot include a backslash () or an @ sign; user names with these characters are treated as Windows names. The user identifier can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.
volume	volume	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage shelves

Storage shelves endpoint overview

Retrieving storage shelf information

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

Examples

1) Retrieve a list of shelves from the cluster

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

```
# The API:
```



```
/api/storage/shelves
```

```
# The call:
```

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

```
# The response:
```

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

```

    }
  }
},
"num_records": 5,
"_links": {
  "self": {
    "href": "/api/storage/shelves/"
  }
}
}
}

```

2) Retrieve a specific shelf from the cluster

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

```

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

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

# The response:
{
  "uid": "3109174803597886800",
  "name": "6.10",
  "id": "10",
  "serial_number": "SHU0954292N0HAH",
  "model": "DS4246",
  "module_type": "iom6",
  "internal": false,
  "local": true,
  "manufacturer": {
    "name": "NETAPP"
  },
  "state": "ok",
  "connection_type": "sas",
  "disk_count": 24,
  "location_led": "off",
  "paths": [
    {

```

```

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

```

```
"id": 1,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 2,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 3,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 4,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 5,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 6,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 7,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 8,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
```

```
},
{
  "id": 9,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 10,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 11,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 12,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 13,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 14,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 15,
  "has_disk": true,
  "type": "single_disk",
  "state": "ok"
},
{
  "id": 16,
  "has_disk": true,
```

```
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 17,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 18,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 19,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 20,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 21,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 22,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 23,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  }
],
```

```
"frus": [
  {
    "type": "module",
    "id": 0,
    "state": "ok",
    "part_number": "111-00690+B2",
    "serial_number": "8001900099",
    "firmware_version": "0191",
    "installed": true
  },
  {
    "type": "module",
    "id": 1,
    "state": "ok",
    "part_number": "111-00190+B0",
    "serial_number": "7903785183",
    "firmware_version": "0191",
    "installed": true
  },
  {
    "type": "psu",
    "id": 1,
    "state": "ok",
    "part_number": "0082562-12",
    "serial_number": "PMW82562007513E",
    "firmware_version": "0311",
    "installed": true,
    "psu": {
      "model": "9C"
    }
  },
  {
    "type": "psu",
    "id": 2,
    "state": "ok",
    "part_number": "0082562-12",
    "serial_number": "PMW825620075138",
    "firmware_version": "0311",
    "installed": true,
    "psu": {
      "model": "9C"
    }
  },
  {
    "type": "psu",
    "id": 3,
```

```

    "state": "ok",
    "part_number": "0082562-12",
    "serial_number": "PMW8256200750BA",
    "firmware_version": "0311",
    "installed": true,
    "psu": {
      "model": "9C"
    }
  },
  {
    "type": "psu",
    "id": 4,
    "state": "ok",
    "part_number": "0082562-12",
    "serial_number": "PMW8256200750A2",
    "firmware_version": "0311",
    "installed": true,
    "psu": {
      "model": "9C"
    }
  }
],
"ports": [
  {
    "id": 0,
    "module_id": "a",
    "designator": "square",
    "state": "connected",
    "internal": false,
    "wwn": "500A098000C9EDBF",
    "cable": {
      "identifier": "5001086000702488-500a098000c9edbf",
      "part_number": "112-00430+A0",
      "length": "2m",
      "serial_number": "APF16510229807"
    },
    "remote": {
      "wwn": "5001086000702488",
      "phy": "08"
    }
  },
  {
    "id": 1,
    "module_id": "a",
    "designator": "circle",
    "state": "connected",

```



```

"internal": false,
"wwn": "500A098000C9EDBF",
"cable": {
  "identifier": "500a098000d5c4bf-500a098000c9edbf",
  "part_number": "112-00176+A0",
  "length": "0.5-1.0m",
  "serial_number": "APF133917610YT"
},
"remote": {
  "wwn": "500A098000D5C4BF",
  "phy": "00"
}
},
{
  "id": 2,
  "module_id": "b",
  "designator": "square",
  "state": "connected",
  "internal": false,
  "wwn": "500A098004F208BF",
  "cable": {
    "identifier": "5001086000702648-500a098004f208bf",
    "part_number": "112-00430+A0",
    "length": "2m",
    "serial_number": "APF16510229540"
  },
  "remote": {
    "wwn": "5001086000702648",
    "phy": "08"
  }
},
{
  "id": 3,
  "module_id": "b",
  "designator": "circle",
  "state": "connected",
  "internal": false,
  "wwn": "500A098004F208BF",
  "cable": {
    "identifier": "500a0980062ba33f-500a098004f208bf",
    "part_number": "112-00176+20",
    "length": "0.5-1.0m",
    "serial_number": "832210017"
  },
  "remote": {
    "wwn": "500A0980062BA33F",

```

```
    "phy": "00"
  }
}
],
"fans": [
  {
    "id": 1,
    "location": "rear of the shelf on the upper left power supply",
    "rpm": 3150,
    "state": "ok",
    "installed": true
  },
  {
    "id": 2,
    "location": "rear of the shelf on the upper left power supply",
    "rpm": 3000,
    "state": "ok",
    "installed": true
  },
  {
    "id": 3,
    "location": "rear of the shelf on the upper right power supply",
    "rpm": 3220,
    "state": "ok",
    "installed": true
  },
  {
    "id": 4,
    "location": "rear of the shelf on the upper right power supply",
    "rpm": 3000,
    "state": "ok",
    "installed": true
  },
  {
    "id": 5,
    "location": "rear of the shelf on the lower left power supply",
    "rpm": 3000,
    "state": "ok",
    "installed": true
  },
  {
    "id": 6,
    "location": "rear of the shelf on the lower left power supply",
    "rpm": 3150,
    "state": "ok",
    "installed": true
  }
]
```

```

},
{
  "id": 7,
  "location": "rear of the shelf on the lower right power supply",
  "rpm": 3150,
  "state": "ok",
  "installed": true
},
{
  "id": 8,
  "location": "rear of the shelf on the lower right power supply",
  "rpm": 3000,
  "state": "ok",
  "installed": true
}
],
"temperature_sensors": [
  {
    "id": 1,
    "location": "front of the shelf on the left, on the OPS panel",
    "temperature": 20,
    "ambient": true,
    "state": "ok",
    "installed": true,
    "threshold": {
      "high": {
        "critical": 42,
        "warning": 40
      },
      "low": {
        "critical": 0,
        "warning": 5
      }
    }
  }
],
{
  "id": 2,
  "location": "inside of the shelf on the midplane",
  "temperature": 29,
  "ambient": false,
  "state": "ok",
  "installed": true,
  "threshold": {
    "high": {
      "critical": 55,
      "warning": 50
    }
  }
}

```

```

    },
    "low": {
      "critical": 5,
      "warning": 10
    }
  }
},
{
  "id": 3,
  "location": "rear of the shelf on the upper left power supply",
  "temperature": 33,
  "ambient": false,
  "state": "ok",
  "installed": true,
  "threshold": {
    "high": {
      "critical": 55,
      "warning": 50
    },
    "low": {
      "critical": 5,
      "warning": 10
    }
  }
},
{
  "id": 4,
  "location": "rear of the shelf on the upper left power supply",
  "temperature": 41,
  "ambient": false,
  "state": "ok",
  "installed": true,
  "threshold": {
    "high": {
      "critical": 70,
      "warning": 65
    },
    "low": {
      "critical": 5,
      "warning": 10
    }
  }
},
{
  "id": 5,
  "location": "rear of the shelf on the upper right power supply",

```

```
"temperature": 32,
"ambient": false,
"state": "ok",
"installed": true,
"threshold": {
  "high": {
    "critical": 55,
    "warning": 50
  },
  "low": {
    "critical": 5,
    "warning": 10
  }
}
},
{
  "id": 6,
  "location": "rear of the shelf on the upper right power supply",
  "temperature": 41,
  "ambient": false,
  "state": "ok",
  "installed": true,
  "threshold": {
    "high": {
      "critical": 70,
      "warning": 65
    },
    "low": {
      "critical": 5,
      "warning": 10
    }
  }
}
},
{
  "id": 7,
  "location": "rear of the shelf on the lower left power supply",
  "temperature": 34,
  "ambient": false,
  "state": "ok",
  "installed": true,
  "threshold": {
    "high": {
      "critical": 55,
      "warning": 50
    },
    "low": {
```

```

        "critical": 5,
        "warning": 10
    }
}
},
{
    "id": 8,
    "location": "rear of the shelf on the lower left power supply",
    "temperature": 45,
    "ambient": false,
    "state": "ok",
    "installed": true,
    "threshold": {
        "high": {
            "critical": 70,
            "warning": 65
        },
        "low": {
            "critical": 5,
            "warning": 10
        }
    }
},
{
    "id": 9,
    "location": "rear of the shelf on the lower right power supply",
    "temperature": 30,
    "ambient": false,
    "state": "ok",
    "installed": true,
    "threshold": {
        "high": {
            "critical": 55,
            "warning": 50
        },
        "low": {
            "critical": 5,
            "warning": 10
        }
    }
},
{
    "id": 10,
    "location": "rear of the shelf on the lower right power supply",
    "temperature": 40,
    "ambient": false,

```

```
"state": "ok",
"installed": true,
"threshold": {
  "high": {
    "critical": 70,
    "warning": 65
  },
  "low": {
    "critical": 5,
    "warning": 10
  }
}
},
{
  "id": 11,
  "location": "rear of the shelf at the top left, on shelf module A",
  "temperature": 30,
  "ambient": false,
  "state": "ok",
  "installed": true,
  "threshold": {
    "high": {
      "critical": 60,
      "warning": 55
    },
    "low": {
      "critical": 5,
      "warning": 10
    }
  }
}
},
{
  "id": 12,
  "location": "rear of the shelf at the top right, on shelf module B",
  "temperature": 33,
  "ambient": false,
  "state": "ok",
  "installed": true,
  "threshold": {
    "high": {
      "critical": 60,
      "warning": 55
    },
    "low": {
      "critical": 5,
      "warning": 10
    }
  }
}
```

```

    }
  }
},
"voltage_sensors": [
  {
    "id": 1,
    "location": "rear of the shelf on the upper left power supply",
    "voltage": 5.11,
    "state": "ok",
    "installed": true
  },
  {
    "id": 2,
    "location": "rear of the shelf on the upper left power supply",
    "voltage": 12.38,
    "state": "ok",
    "installed": true
  },
  {
    "id": 3,
    "location": "rear of the shelf on the upper right power supply",
    "voltage": 5.11,
    "state": "ok",
    "installed": true
  },
  {
    "id": 4,
    "location": "rear of the shelf on the upper right power supply",
    "voltage": 12.26,
    "state": "ok",
    "installed": true
  },
  {
    "id": 5,
    "location": "rear of the shelf on the lower left power supply",
    "voltage": 5.7,
    "state": "ok",
    "installed": true
  },
  {
    "id": 6,
    "location": "rear of the shelf on the lower left power supply",
    "voltage": 12.26,
    "state": "ok",
    "installed": true
  }
]

```



```
},
{
  "id": 7,
  "location": "rear of the shelf on the lower right power supply",
  "voltage": 5.15,
  "state": "ok",
  "installed": true
},
{
  "id": 8,
  "location": "rear of the shelf on the lower right power supply",
  "voltage": 12.3,
  "state": "ok",
  "installed": true
}
],
"current_sensors": [
  {
    "id": 1,
    "location": "rear of the shelf on the upper left power supply",
    "current": 6990,
    "state": "ok",
    "installed": true
  },
  {
    "id": 2,
    "location": "rear of the shelf on the upper left power supply",
    "current": 5150,
    "state": "ok",
    "installed": true
  },
  {
    "id": 3,
    "location": "rear of the shelf on the upper right power supply",
    "current": 4600,
    "state": "ok",
    "installed": true
  },
  {
    "id": 4,
    "location": "rear of the shelf on the upper right power supply",
    "current": 4800,
    "state": "ok",
    "installed": true
  },
  {
```

```

    "id": 5,
    "location": "rear of the shelf on the lower left power supply",
    "current": 4140,
    "state": "ok",
    "installed": true
  },
  {
    "id": 6,
    "location": "rear of the shelf on the lower left power supply",
    "current": 7770,
    "state": "ok",
    "installed": true
  },
  {
    "id": 7,
    "location": "rear of the shelf on the lower right power supply",
    "current": 4140,
    "state": "ok",
    "installed": true
  },
  {
    "id": 8,
    "location": "rear of the shelf on the lower right power supply",
    "current": 4720,
    "state": "ok",
    "installed": true
  }
],
"acps": [
  {
    "enabled": true,
    "channel": "in_band",
    "connection_state": "active",
    "node": {
      "uuid": "cf62d23c-6100-11eb-9852-00a098fd725d",
      "name": "cat33-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/cf62d23c-6100-11eb-9852-00a098fd725d"
        }
      }
    }
  },
  {
    "enabled": true,

```

```

    "channel": "in_band",
    "connection_state": "active",
    "node": {
      "uuid": "d0892dd7-6100-11eb-9cdb-d039ea010238",
      "name": "cat33-02",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/d0892dd7-6100-11eb-9cdb-
d039ea010238"
        }
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/storage/shelves/3109174803597886800"
    }
  }
}

```

Modifying storage shelf

The storage shelf PATCH API modifies the shelf location LED.

Example

```

# The API:
/api/storage/shelves

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/shelves/3109174803597886800"
-H "accept: application/hal+json" -H "Content-Type: application/hal+json"
-d '{"location_led": "on"}'

# The response:
{
}

```

Retrieve shelves

GET /storage/shelves

Introduced In: 9.6

Retrieves a collection of shelves.

Related ONTAP commands

- `storage shelf show`
- `storage shelf port show`
- `storage shelf drawer show`
- `storage shelf drawer show-slot`
- `storage shelf acp show`

Learn more

- [DOC /storage/shelves](#)

Parameters

Name	Type	In	Required	Description
acps.channel	string	query	False	Filter by acps.channel <ul style="list-style-type: none">• Introduced in: 9.10
acps.connection_state	string	query	False	Filter by acps.connection_state <ul style="list-style-type: none">• Introduced in: 9.10
acps.port	string	query	False	Filter by acps.port <ul style="list-style-type: none">• Introduced in: 9.10
acps.error.severity	string	query	False	Filter by acps.error.severity <ul style="list-style-type: none">• Introduced in: 9.10

Name	Type	In	Required	Description
acps.error.reason.arguments.message	string	query	False	Filter by acps.error.reason.arguments.message • Introduced in: 9.10
acps.error.reason.arguments.code	string	query	False	Filter by acps.error.reason.arguments.code • Introduced in: 9.10
acps.error.reason.message	string	query	False	Filter by acps.error.reason.message • Introduced in: 9.10
acps.error.reason.code	string	query	False	Filter by acps.error.reason.code • Introduced in: 9.10
acps.error.reason.target	string	query	False	Filter by acps.error.reason.target • Introduced in: 9.10
acps.error.type	string	query	False	Filter by acps.error.type • Introduced in: 9.10
acps.node.uuid	string	query	False	Filter by acps.node.uuid • Introduced in: 9.10

Name	Type	In	Required	Description
acps.node.name	string	query	False	Filter by acps.node.name • Introduced in: 9.10
acps.subnet	string	query	False	Filter by acps.subnet • Introduced in: 9.10
acps.address	string	query	False	Filter by acps.address • Introduced in: 9.10
acps.enabled	boolean	query	False	Filter by acps.enabled • Introduced in: 9.10
acps.netmask	string	query	False	Filter by acps.netmask • Introduced in: 9.10
id	string	query	False	Filter by id
uid	string	query	False	Filter by uid
connection_type	string	query	False	Filter by connection_type
vendor.product	string	query	False	Filter by vendor.product • Introduced in: 9.8
vendor.name	string	query	False	Filter by vendor.name • Introduced in: 9.10

Name	Type	In	Required	Description
vendor.part_number	string	query	False	Filter by vendor.part_number • Introduced in: 9.8
vendor.serial_number	string	query	False	Filter by vendor.serial_number • Introduced in: 9.8
vendor.manufacturer	string	query	False	Filter by vendor.manufacturer • Introduced in: 9.8
bays.state	string	query	False	Filter by bays.state
bays.id	integer	query	False	Filter by bays.id
bays.type	string	query	False	Filter by bays.type
bays.has_disk	boolean	query	False	Filter by bays.has_disk
bays.drawer.id	integer	query	False	Filter by bays.drawer.id • Introduced in: 9.11
bays.drawer.slot	integer	query	False	Filter by bays.drawer.slot • Introduced in: 9.11
location_led	string	query	False	Filter by location_led • Introduced in: 9.10

Name	Type	In	Required	Description
current_sensors.current	integer	query	False	Filter by current_sensors.current • Introduced in: 9.10
current_sensors.id	integer	query	False	Filter by current_sensors.id • Introduced in: 9.10
current_sensors.location	string	query	False	Filter by current_sensors.location • Introduced in: 9.10
current_sensors.state	string	query	False	Filter by current_sensors.state • Introduced in: 9.10
current_sensors.installed	boolean	query	False	Filter by current_sensors.installed • Introduced in: 9.13
errors.reason.arguments.message	string	query	False	Filter by errors.reason.arguments.message • Introduced in: 9.10
errors.reason.arguments.code	string	query	False	Filter by errors.reason.arguments.code • Introduced in: 9.10

Name	Type	In	Required	Description
errors.reason.message	string	query	False	Filter by errors.reason.message • Introduced in: 9.9
errors.reason.code	string	query	False	Filter by errors.reason.code • Introduced in: 9.9
errors.reason.target	string	query	False	Filter by errors.reason.target • Introduced in: 9.10
serial_number	string	query	False	Filter by serial_number
manufacturer.name	string	query	False	Filter by manufacturer.name • Introduced in: 9.10
model	string	query	False	Filter by model
voltage_sensors.id	integer	query	False	Filter by voltage_sensors.id • Introduced in: 9.10
voltage_sensors.location	string	query	False	Filter by voltage_sensors.location • Introduced in: 9.10
voltage_sensors.installed	boolean	query	False	Filter by voltage_sensors.installed • Introduced in: 9.13

Name	Type	In	Required	Description
voltage_sensors.state	string	query	False	Filter by voltage_sensors.state • Introduced in: 9.10
voltage_sensors.voltage	number	query	False	Filter by voltage_sensors.voltage • Introduced in: 9.10
drawers.error	string	query	False	Filter by drawers.error
drawers.state	string	query	False	Filter by drawers.state
drawers.closed	boolean	query	False	Filter by drawers.closed
drawers.part_number	string	query	False	Filter by drawers.part_number
drawers.serial_number	string	query	False	Filter by drawers.serial_number
drawers.id	integer	query	False	Filter by drawers.id
drawers.disk_count	integer	query	False	Filter by drawers.disk_count
local	boolean	query	False	Filter by local • Introduced in: 9.8
internal	boolean	query	False	Filter by internal
ports.id	integer	query	False	Filter by ports.id
ports.internal	boolean	query	False	Filter by ports.internal

Name	Type	In	Required	Description
ports.module_id	string	query	False	Filter by ports.module_id
ports.state	string	query	False	Filter by ports.state
ports.mac_address	string	query	False	Filter by ports.mac_address
ports.cable.identifier	string	query	False	Filter by ports.cable.identifier
ports.cable.length	string	query	False	Filter by ports.cable.length
ports.cable.part_number	string	query	False	Filter by ports.cable.part_number
ports.cable.serial_number	string	query	False	Filter by ports.cable.serial_number
ports.wwn	string	query	False	Filter by ports.wwn
ports.remote.device	string	query	False	Filter by ports.remote.device • Introduced in: 9.8
ports.remote.phy	string	query	False	Filter by ports.remote.phy
ports.remote.wwn	string	query	False	Filter by ports.remote.wwn
ports.remote.port	string	query	False	Filter by ports.remote.port
ports.remote.mac_address	string	query	False	Filter by ports.remote.mac_address
ports.remote.chassis	string	query	False	Filter by ports.remote.chassis

Name	Type	In	Required	Description
ports.designator	string	query	False	Filter by ports.designator
frus.psu.power_rating	integer	query	False	Filter by frus.psu.power_rating • Introduced in: 9.10
frus.psu.model	string	query	False	Filter by frus.psu.model • Introduced in: 9.10
frus.psu.power_drawn	integer	query	False	Filter by frus.psu.power_drawn • Introduced in: 9.10
frus.psu.crest_factor	integer	query	False	Filter by frus.psu.crest_factor • Introduced in: 9.10
frus.type	string	query	False	Filter by frus.type
frus.id	integer	query	False	Filter by frus.id
frus.firmware_version	string	query	False	Filter by frus.firmware_version
frus.state	string	query	False	Filter by frus.state
frus.installed	boolean	query	False	Filter by frus.installed • Introduced in: 9.10
frus.serial_number	string	query	False	Filter by frus.serial_number

Name	Type	In	Required	Description
frus.part_number	string	query	False	Filter by frus.part_number
fans.state	string	query	False	Filter by fans.state <ul style="list-style-type: none"> • Introduced in: 9.9
fans.installed	boolean	query	False	Filter by fans.installed <ul style="list-style-type: none"> • Introduced in: 9.13
fans.location	string	query	False	Filter by fans.location <ul style="list-style-type: none"> • Introduced in: 9.9
fans.rpm	integer	query	False	Filter by fans.rpm <ul style="list-style-type: none"> • Introduced in: 9.9
fans.id	integer	query	False	Filter by fans.id <ul style="list-style-type: none"> • Introduced in: 9.9
disk_count	integer	query	False	Filter by disk_count
state	string	query	False	Filter by state
paths.name	string	query	False	Filter by paths.name
paths.node.uuid	string	query	False	Filter by paths.node.uuid
paths.node.name	string	query	False	Filter by paths.node.name
name	string	query	False	Filter by name

Name	Type	In	Required	Description
temperature_sensors.threshold.low.critical	integer	query	False	Filter by temperature_sensors.threshold.low.critical • Introduced in: 9.10
temperature_sensors.threshold.low.warning	integer	query	False	Filter by temperature_sensors.threshold.low.warning • Introduced in: 9.10
temperature_sensors.threshold.high.warning	integer	query	False	Filter by temperature_sensors.threshold.high.warning • Introduced in: 9.10
temperature_sensors.threshold.high.critical	integer	query	False	Filter by temperature_sensors.threshold.high.critical • Introduced in: 9.10
temperature_sensors.installed	boolean	query	False	Filter by temperature_sensors.installed • Introduced in: 9.13
temperature_sensors.temperature	integer	query	False	Filter by temperature_sensors.temperature • Introduced in: 9.10

Name	Type	In	Required	Description
temperature_sensor s.state	string	query	False	Filter by temperature_sensor s.state • Introduced in: 9.10
temperature_sensor s.location	string	query	False	Filter by temperature_sensor s.location • Introduced in: 9.10
temperature_sensor s.id	integer	query	False	Filter by temperature_sensor s.id • Introduced in: 9.10
temperature_sensor s.ambient	boolean	query	False	Filter by temperature_sensor s.ambient • Introduced in: 9.10
module_type	string	query	False	Filter by module_type
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • 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[shelf]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "acps": {
      "address": "192.168.1.104",
      "channel": "out_of_band",
      "connection_state": "full_connectivity",
      "error": {
        "reason": {
          "arguments": {
            "code": "string",
            "message": "string"
          },
          "code": "4",
          "message": "entry doesn't exist",
          "target": "uuid"
        },
        "severity": "unknown",
        "type": "not_applicable"
      },
      "netmask": "255.255.252.0",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "port": "e0P",
      "subnet": "192.168.0.1"
    },
    "bays": {
      "drawer": {
        "id": 1,

```

```

    "slot": 0
  },
  "id": 0,
  "state": "ok",
  "type": "single_disk"
},
"connection_type": "sas",
"current_sensors": {
  "current": 14410,
  "id": 1,
  "installed": 1,
  "location": "rear of the shelf on the lower left power supply",
  "state": "ok"
},
"disk_count": 12,
"drawers": {
  "disk_count": 12,
  "part_number": "111-03071",
  "serial_number": "021604008263",
  "state": "ok"
},
"errors": {
  "reason": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
},
"fans": {
  "id": 1,
  "installed": 1,
  "location": "rear of the shelf on the lower left power supply",
  "rpm": 3020,
  "state": "ok"
},
"frus": {
  "firmware_version": "0191",
  "installed": 1,
  "part_number": "111-00690+A2",
  "psu": {
    "crest_factor": 92,
    "model": "00",

```

```

    "power_drawn": 210,
    "power_rating": 1600
  },
  "serial_number": "8000166294",
  "state": "error",
  "type": "module"
},
"id": "1",
"location_led": "off",
"manufacturer": {
  "name": "NETAPP"
},
"model": "DS2246",
"module_type": "iom6",
"name": "1.1",
"paths": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "2a",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"ports": {
  "cable": {
    "identifier": "500a098000b6c3f-50000d1703544b80",
    "length": "2m",
    "part_number": "112-00431+A0",
    "serial_number": "616930439"
  },
  "designator": "square",
  "id": 0,
  "module_id": "a",
  "remote": {
    "phy": "12",
    "wwn": "50000D1703544B80"
  }
},

```

```
    "state": "connected",
    "wwn": "500A0980000B6C3F"
  },
  "serial_number": "SHFMS1514000895",
  "state": "ok",
  "temperature_sensors": {
    "ambient": "",
    "id": 1,
    "installed": 1,
    "location": "temp sensor on midplane left",
    "state": "ok",
    "temperature": 32,
    "threshold": {
      "high": {
        "critical": 60,
        "warning": 55
      },
      "low": {
        "critical": 0,
        "warning": 5
      }
    }
  },
  "uid": "7777841915827391056",
  "vendor": {
    "manufacturer": "XYZ",
    "name": "XYZ",
    "part_number": "A92831142733",
    "product": "LS2246",
    "serial_number": "891234572210221"
  },
  "voltage_sensors": {
    "id": 1,
    "installed": 1,
    "location": "rear of the shelf on the lower left power supply",
    "state": "ok",
    "voltage": 12.18
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

error

Error object is populated when connection_state becomes non-optimal

Name	Type	Description
reason	error	
severity	string	
type	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

acps

Name	Type	Description
address	string	
channel	string	
connection_state	string	
enabled	boolean	
error	error	Error object is populated when connection_state becomes non-optimal
netmask	string	
node	node	
port	string	
subnet	string	

drawer

Name	Type	Description
id	integer	The drawer containing this bay
slot	integer	The drawer slot for this bay

bays

Name	Type	Description
drawer	drawer	
has_disk	boolean	
id	integer	
state	string	

Name	Type	Description
type	string	

current_sensors

Name	Type	Description
current	integer	Current, in milliamps
id	integer	
installed	boolean	
location	string	
state	string	

drawers

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

errors

Name	Type	Description
reason	error	

fans

Name	Type	Description
id	integer	
installed	boolean	
location	string	
rpm	integer	
state	string	

psu

Name	Type	Description
crest_factor	integer	The ratio of the peak voltage to the root-mean-square voltage
model	string	
power_drawn	integer	Power drawn, in watts
power_rating	integer	Power rating, in watts

frus

Name	Type	Description
firmware_version	string	
id	integer	
installed	boolean	
part_number	string	
psu	psu	
serial_number	string	
state	string	
type	string	

manufacturer

Name	Type	Description
name	string	

paths

Storage port

Name	Type	Description
_links	_links	
name	string	
node	node	

cable

Name	Type	Description
identifier	string	
length	string	

Name	Type	Description
part_number	string	
serial_number	string	

remote

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

ports

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

high

Name	Type	Description
critical	integer	High critical threshold, in degrees Celsius
warning	integer	High warning threshold, in degrees Celsius

low

Name	Type	Description
critical	integer	Low critical threshold, in degrees Celsius
warning	integer	Low warning threshold, in degrees Celsius

threshold

Name	Type	Description
high	high	
low	low	

temperature_sensors

Name	Type	Description
ambient	boolean	Sensor that measures the ambient temperature
id	integer	
installed	boolean	
location	string	
state	string	
temperature	integer	Temperature, in degrees Celsius
threshold	threshold	

vendor

Name	Type	Description
manufacturer	string	Support for this field will be removed in a future release. Please use vendor.name for this field.
name	string	
part_number	string	Part number
product	string	Product name
serial_number	string	Serial number

voltage_sensors

Name	Type	Description
id	integer	
installed	boolean	
location	string	
state	string	
voltage	number	Voltage, in volts

shelf

Name	Type	Description
acps	array[acps]	Alternate Control Paths to ACP processors/functions in shelf modules and expanders
bays	array[bays]	
connection_type	string	
current_sensors	array[current_sensors]	
disk_count	integer	
drawers	array[drawers]	
errors	array[errors]	
fans	array[fans]	
frus	array[frus]	
id	string	
internal	boolean	
local	boolean	
location_led	string	
manufacturer	manufacturer	
model	string	
module_type	string	
name	string	
paths	array[paths]	
ports	array[ports]	
serial_number	string	
state	string	
temperature_sensors	array[temperature_sensors]	
uid	string	

Name	Type	Description
vendor	vendor	
voltage_sensors	array[voltage_sensors]	

Retrieve a shelf

GET /storage/shelves/{uid}

Introduced In: 9.6

Retrieves a specific shelf.

Related ONTAP commands

- storage shelf show
- storage shelf port show
- storage shelf drawer show
- storage shelf drawer show-slot
- storage shelf acp show

Learn more

- [DOC /storage/shelves](#)

Parameters

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

Response

Status: 200, Ok

Name	Type	Description
acps	array[acps]	Alternate Control Paths to ACP processors/functions in shelf modules and expanders
bays	array[bays]	

Name	Type	Description
connection_type	string	
current_sensors	array[current_sensors]	
disk_count	integer	
drawers	array[drawers]	
errors	array[errors]	
fans	array[fans]	
frus	array[frus]	
id	string	
internal	boolean	
local	boolean	
location_led	string	
manufacturer	manufacturer	
model	string	
module_type	string	
name	string	
paths	array[paths]	
ports	array[ports]	
serial_number	string	
state	string	
temperature_sensors	array[temperature_sensors]	
uid	string	
vendor	vendor	
voltage_sensors	array[voltage_sensors]	

Example response

```
{
  "acps": {
    "address": "192.168.1.104",
    "channel": "out_of_band",
    "connection_state": "full_connectivity",
    "error": {
      "reason": {
        "arguments": {
          "code": "string",
          "message": "string"
        },
        "code": "4",
        "message": "entry doesn't exist",
        "target": "uuid"
      },
      "severity": "unknown",
      "type": "not_applicable"
    },
    "netmask": "255.255.252.0",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "port": "e0P",
    "subnet": "192.168.0.1"
  },
  "bays": {
    "drawer": {
      "id": 1,
      "slot": 0
    },
    "id": 0,
    "state": "ok",
    "type": "single_disk"
  },
  "connection_type": "sas",
  "current_sensors": {
    "current": 14410,
    "id": 1,
```

```

    "installed": 1,
    "location": "rear of the shelf on the lower left power supply",
    "state": "ok"
  },
  "disk_count": 12,
  "drawers": {
    "disk_count": 12,
    "part_number": "111-03071",
    "serial_number": "021604008263",
    "state": "ok"
  },
  "errors": {
    "reason": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    }
  },
  "fans": {
    "id": 1,
    "installed": 1,
    "location": "rear of the shelf on the lower left power supply",
    "rpm": 3020,
    "state": "ok"
  },
  "frus": {
    "firmware_version": "0191",
    "installed": 1,
    "part_number": "111-00690+A2",
    "psu": {
      "crest_factor": 92,
      "model": "00",
      "power_drawn": 210,
      "power_rating": 1600
    },
    "serial_number": "8000166294",
    "state": "error",
    "type": "module"
  },
  "id": "1",
  "location_led": "off",
  "manufacturer": {

```



```

    "name": "NETAPP"
  },
  "model": "DS2246",
  "module_type": "iom6",
  "name": "1.1",
  "paths": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "2a",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "ports": {
    "cable": {
      "identifier": "500a0980000b6c3f-50000d1703544b80",
      "length": "2m",
      "part_number": "112-00431+A0",
      "serial_number": "616930439"
    },
    "designator": "square",
    "id": 0,
    "module_id": "a",
    "remote": {
      "phy": "12",
      "wwn": "50000D1703544B80"
    },
    "state": "connected",
    "wwn": "500A0980000B6C3F"
  },
  "serial_number": "SHFMS1514000895",
  "state": "ok",
  "temperature_sensors": {
    "ambient": "",
    "id": 1,
    "installed": 1,
    "location": "temp sensor on midplane left",

```

```

"state": "ok",
"temperature": 32,
"threshold": {
  "high": {
    "critical": 60,
    "warning": 55
  },
  "low": {
    "critical": 0,
    "warning": 5
  }
},
"uid": "7777841915827391056",
"vendor": {
  "manufacturer": "XYZ",
  "name": "XYZ",
  "part_number": "A92831142733",
  "product": "LS2246",
  "serial_number": "891234572210221"
},
"voltage_sensors": {
  "id": 1,
  "installed": 1,
  "location": "rear of the shelf on the lower left power supply",
  "state": "ok",
  "voltage": 12.18
}
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

error

Error object is populated when connection_state becomes non-optimal

Name	Type	Description
reason	error	
severity	string	
type	string	

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	
uuid	string	

acps

Name	Type	Description
address	string	
channel	string	
connection_state	string	
enabled	boolean	
error	error	Error object is populated when connection_state becomes non-optimal
netmask	string	
node	node	
port	string	
subnet	string	

drawer

Name	Type	Description
id	integer	The drawer containing this bay
slot	integer	The drawer slot for this bay

bays

Name	Type	Description
drawer	drawer	
has_disk	boolean	
id	integer	
state	string	
type	string	

current_sensors

Name	Type	Description
current	integer	Current, in milliamps

Name	Type	Description
id	integer	
installed	boolean	
location	string	
state	string	

drawers

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

errors

Name	Type	Description
reason	error	

fans

Name	Type	Description
id	integer	
installed	boolean	
location	string	
rpm	integer	
state	string	

psu

Name	Type	Description
crest_factor	integer	The ratio of the peak voltage to the root-mean-square voltage
model	string	
power_drawn	integer	Power drawn, in watts

Name	Type	Description
power_rating	integer	Power rating, in watts

frus

Name	Type	Description
firmware_version	string	
id	integer	
installed	boolean	
part_number	string	
psu	psu	
serial_number	string	
state	string	
type	string	

manufacturer

Name	Type	Description
name	string	

paths

Storage port

Name	Type	Description
_links	_links	
name	string	
node	node	

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

remote

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

ports

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

high

Name	Type	Description
critical	integer	High critical threshold, in degrees Celsius
warning	integer	High warning threshold, in degrees Celsius

low

Name	Type	Description
critical	integer	Low critical threshold, in degrees Celsius
warning	integer	Low warning threshold, in degrees Celsius

threshold

Name	Type	Description
high	high	
low	low	

temperature_sensors

Name	Type	Description
ambient	boolean	Sensor that measures the ambient temperature
id	integer	
installed	boolean	
location	string	
state	string	
temperature	integer	Temperature, in degrees Celsius
threshold	threshold	

vendor

Name	Type	Description
manufacturer	string	Support for this field will be removed in a future release. Please use vendor.name for this field.
name	string	
part_number	string	Part number
product	string	Product name
serial_number	string	Serial number

voltage_sensors

Name	Type	Description
id	integer	
installed	boolean	
location	string	
state	string	
voltage	number	Voltage, in volts

Update a shelf location LED

PATCH /storage/shelves/{uid}

Introduced In: 9.10

Updates a shelf location LED.

Related ONTAP commands

- `storage shelf location-led modify`

Learn more

- [DOC /storage/shelves](#)

Parameters

Name	Type	In	Required	Description
uid	string	path	True	Shelf UID

Request Body

Name	Type	Description
acps	array[acps]	Alternate Control Paths to ACP processors/functions in shelf modules and expanders
bays	array[bays]	
connection_type	string	
current_sensors	array[current_sensors]	
disk_count	integer	
drawers	array[drawers]	
errors	array[errors]	
fans	array[fans]	
frus	array[frus]	
id	string	
internal	boolean	
local	boolean	
location_led	string	
manufacturer	manufacturer	
model	string	

Name	Type	Description
module_type	string	
name	string	
paths	array[paths]	
ports	array[ports]	
serial_number	string	
state	string	
temperature_sensors	array[temperature_sensors]	
uid	string	
vendor	vendor	
voltage_sensors	array[voltage_sensors]	

Example request

```
{
  "acps": {
    "address": "192.168.1.104",
    "channel": "out_of_band",
    "connection_state": "full_connectivity",
    "error": {
      "reason": {
        "arguments": {
          "code": "string",
          "message": "string"
        },
        "code": "4",
        "message": "entry doesn't exist",
        "target": "uuid"
      },
      "severity": "unknown",
      "type": "not_applicable"
    },
    "netmask": "255.255.252.0",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "port": "e0P",
    "subnet": "192.168.0.1"
  },
  "bays": {
    "drawer": {
      "id": 1,
      "slot": 0
    },
    "id": 0,
    "state": "ok",
    "type": "single_disk"
  },
  "connection_type": "sas",
  "current_sensors": {
    "current": 14410,
    "id": 1,
```

```

    "installed": 1,
    "location": "rear of the shelf on the lower left power supply",
    "state": "ok"
  },
  "disk_count": 12,
  "drawers": {
    "disk_count": 12,
    "part_number": "111-03071",
    "serial_number": "021604008263",
    "state": "ok"
  },
  "errors": {
    "reason": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    }
  },
  "fans": {
    "id": 1,
    "installed": 1,
    "location": "rear of the shelf on the lower left power supply",
    "rpm": 3020,
    "state": "ok"
  },
  "frus": {
    "firmware_version": "0191",
    "installed": 1,
    "part_number": "111-00690+A2",
    "psu": {
      "crest_factor": 92,
      "model": "00",
      "power_drawn": 210,
      "power_rating": 1600
    },
    "serial_number": "8000166294",
    "state": "error",
    "type": "module"
  },
  "id": "1",
  "location_led": "off",
  "manufacturer": {

```

```

    "name": "NETAPP"
  },
  "model": "DS2246",
  "module_type": "iom6",
  "name": "1.1",
  "paths": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "2a",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "ports": {
    "cable": {
      "identifier": "500a0980000b6c3f-50000d1703544b80",
      "length": "2m",
      "part_number": "112-00431+A0",
      "serial_number": "616930439"
    },
    "designator": "square",
    "id": 0,
    "module_id": "a",
    "remote": {
      "phy": "12",
      "wwn": "50000D1703544B80"
    },
    "state": "connected",
    "wwn": "500A0980000B6C3F"
  },
  "serial_number": "SHFMS1514000895",
  "state": "ok",
  "temperature_sensors": {
    "ambient": "",
    "id": 1,
    "installed": 1,
    "location": "temp sensor on midplane left",

```

```

"state": "ok",
"temperature": 32,
"threshold": {
  "high": {
    "critical": 60,
    "warning": 55
  },
  "low": {
    "critical": 0,
    "warning": 5
  }
},
"uid": "7777841915827391056",
"vendor": {
  "manufacturer": "XYZ",
  "name": "XYZ",
  "part_number": "A92831142733",
  "product": "LS2246",
  "serial_number": "891234572210221"
},
"voltage_sensors": {
  "id": 1,
  "installed": 1,
  "location": "rear of the shelf on the lower left power supply",
  "state": "ok",
  "voltage": 12.18
}
}

```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
17825872	Shelf locate request failed because shelf "<name>" was not found.</name>

Error Code	Description
17825873	Shelf locate request failed because shelf "<name>" does not support this command.</name>
17825874	Shelf locate request failed for shelf "<name>" with an unknown error.</name>
17825875	Shelf locate request failed for shelf "<name>" because shelf modules are unreachable.</name>

Name	Type	Description
error	error	

Example error

```

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

```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

error

Error object is populated when connection_state becomes non-optimal

Name	Type	Description
reason	error	
severity	string	
type	string	

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	
uuid	string	

acps

Name	Type	Description
address	string	
channel	string	
connection_state	string	
enabled	boolean	
error	error	Error object is populated when connection_state becomes non-optimal
netmask	string	
node	node	
port	string	
subnet	string	

drawer

Name	Type	Description
id	integer	The drawer containing this bay
slot	integer	The drawer slot for this bay

bays

Name	Type	Description
drawer	drawer	
has_disk	boolean	
id	integer	
state	string	
type	string	

current_sensors

Name	Type	Description
current	integer	Current, in milliamps

Name	Type	Description
id	integer	
installed	boolean	
location	string	
state	string	

drawers

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

errors

Name	Type	Description
reason	error	

fans

Name	Type	Description
id	integer	
installed	boolean	
location	string	
rpm	integer	
state	string	

psu

Name	Type	Description
crest_factor	integer	The ratio of the peak voltage to the root-mean-square voltage
model	string	
power_drawn	integer	Power drawn, in watts

Name	Type	Description
power_rating	integer	Power rating, in watts

frus

Name	Type	Description
firmware_version	string	
id	integer	
installed	boolean	
part_number	string	
psu	psu	
serial_number	string	
state	string	
type	string	

manufacturer

Name	Type	Description
name	string	

paths

Storage port

Name	Type	Description
_links	_links	
name	string	
node	node	

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

remote

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

ports

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

high

Name	Type	Description
critical	integer	High critical threshold, in degrees Celsius
warning	integer	High warning threshold, in degrees Celsius

low

Name	Type	Description
critical	integer	Low critical threshold, in degrees Celsius
warning	integer	Low warning threshold, in degrees Celsius

threshold

Name	Type	Description
high	high	
low	low	

temperature_sensors

Name	Type	Description
ambient	boolean	Sensor that measures the ambient temperature
id	integer	
installed	boolean	
location	string	
state	string	
temperature	integer	Temperature, in degrees Celsius
threshold	threshold	

vendor

Name	Type	Description
manufacturer	string	Support for this field will be removed in a future release. Please use vendor.name for this field.
name	string	
part_number	string	Part number
product	string	Product name
serial_number	string	Serial number

voltage_sensors

Name	Type	Description
id	integer	
installed	boolean	
location	string	
state	string	
voltage	number	Voltage, in volts

shelf

Name	Type	Description
acps	array[acps]	Alternate Control Paths to ACP processors/functions in shelf modules and expanders
bays	array[bays]	
connection_type	string	
current_sensors	array[current_sensors]	
disk_count	integer	
drawers	array[drawers]	
errors	array[errors]	
fans	array[fans]	
frus	array[frus]	
id	string	
internal	boolean	
local	boolean	
location_led	string	
manufacturer	manufacturer	
model	string	
module_type	string	
name	string	
paths	array[paths]	
ports	array[ports]	
serial_number	string	
state	string	
temperature_sensors	array[temperature_sensors]	
uid	string	
vendor	vendor	
voltage_sensors	array[voltage_sensors]	

Create and retrieve Snapshot copy policies

Storage snapshot-policies endpoint overview

Overview

In ONTAP, scheduled Snapshot copy creation works based on Snapshot copy policies. ONTAP provides three cluster-wide Snapshot copy policies: "default", "default-1weekly" and "none". A Snapshot copy policy can have more than one schedule associated with it. A Snapshot copy policy can be linked to a storage object and based on the schedule in the policy, Snapshot copies will be created on the object at that interval. Each schedule in a Snapshot copy policy has a Snapshot copy name prefix attached to it. Every Snapshot copy created using this policy will have this prefix in its name. There is also a retention count associated with every schedule. This count indicates the maximum number of Snapshot copies that can exist for a given schedule. Once the Snapshot copy count reaches the retention count, on the next create operation, the oldest Snapshot copy is deleted. A retention period can be associated with every schedule. During Snapshot copy creation, this period is set as SnapLock expiry time on Snapshot copy locking enabled volumes.

Snapshot copy policy APIs

The following APIs are used to perform operations related to Snapshot copy policy information:

– POST /api/storage/snapshot-policies

– GET /api/storage/snapshot-policies

– GET /api/storage/snapshot-policies/{uuid}

– PATCH /api/storage/snapshot-policies/{uuid}

– DELETE /api/storage/snapshot-policies/{uuid}

Examples

Creating a Snapshot copy policy

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


```
# The API:
/api/storage/snapshot-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/snapshot-policies" -H
"accept: application/hal+json" -d '{"name": "new_policy", "enabled":
"true", "comment": "policy comment", "copies": [{ "schedule": { "name":
"5min" }, "count": "5", "prefix": "xyz" , "retention_period": "PT20M" }],
"svm": { "name": "vs0"} }'
```

```
# The response:
HTTP/1.1 201 Created
Date: Tue, 12 Mar 2019 21:20:24 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/snapshot-policies/a69d8173-450c-11e9-aa44-
005056bbc848
Content-Length: 369
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "uuid": "a69d8173-450c-11e9-aa44-005056bbc848",
      "svm": {
        "name": "vs0"
      },
      "name": "new_policy",
      "comment": "This is a 5min schedule policy",
      "enabled": true,
      "copies": [
        {
          "count": 5,
          "snapmirror_label": "-",
          "retention_period": "PT20M",
          "schedule": {
            "name": "5min"
          }
        }
      ]
    }
  ]
}
```

Retrieving Snapshot copy policy attributes

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

```
# The API:
/api/storage/snapshot-policies

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

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

```

        "href": "/api/storage/snapshot-policies/3c1c1656-2fe8-11e9-b55e-005056bbf1c8"
      }
    },
    {
      "uuid": "3c228b82-2fe8-11e9-b55e-005056bbf1c8",
      "name": "none",
      "_links": {
        "self": {
          "href": "/api/storage/snapshot-policies/3c228b82-2fe8-11e9-b55e-005056bbf1c8"
        }
      }
    }
  ],
  "num_records": 4,
  "_links": {
    "self": {
      "href": "/api/storage/snapshot-policies/"
    }
  }
}

```

Retrieving the attributes of a specific Snapshot copy policy

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

```

# The API:
/api/storage/snapshot-policies/{uuid}

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

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:24:48 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 381
Content-Type: application/json
{
  "uuid": "3c112527-2fe8-11e9-b55e-005056bbf1c8",
  "name": "default",

```

```
"comment": "Default policy with hourly, daily & weekly schedules.",
"enabled": true,
"scope": "cluster",
"copies": [
  {
    "count": 6,
    "prefix": "hourly",
    "schedule": {
      "name": "hourly"
    }
  },
  {
    "count": 2,
    "prefix": "daily",
    "schedule": {
      "name": "daily"
    }
  },
  {
    "count": 2,
    "prefix": "weekly",
    "schedule": {
      "name": "weekly"
    }
  }
],
"_links": {
  "self": {
    "href": "/api/storage/snapshot-policies/3c112527-2fe8-11e9-b55e-005056bbf1c8"
  }
}
}
```

Updating a Snapshot copy policy

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

```
# The API:
/api/storage/snapshot-policies/{uuid}

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

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

Deleting a Snapshot copy policy

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

```
# The API:
/api/storage/snapshot-policies/{uuid}

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

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

Retrieve Snapshot copy policies

GET /storage/snapshot-policies

Introduced In: 9.6

Retrieves a collection of Snapshot copy policies.

Related ONTAP commands

- `snapshot policy show`

Learn more

- [DOC /storage/snapshot-policies](#)

Parameters

Name	Type	In	Required	Description
<code>copies.prefix</code>	string	query	False	Filter by <code>copies.prefix</code>
<code>copies.snapmirror_label</code>	string	query	False	Filter by <code>copies.snapmirror_label</code>
<code>copies.count</code>	integer	query	False	Filter by <code>copies.count</code>
<code>copies.retention_period</code>	string	query	False	Filter by <code>copies.retention_period</code> <ul style="list-style-type: none">• Introduced in: 9.12
<code>copies.schedule.name</code>	string	query	False	Filter by <code>copies.schedule.name</code>
<code>name</code>	string	query	False	Filter by name
<code>uuid</code>	string	query	False	Filter by uuid
<code>comment</code>	string	query	False	Filter by comment
<code>enabled</code>	boolean	query	False	Filter by enabled
<code>scope</code>	string	query	False	Filter by scope
<code>svm.uuid</code>	string	query	False	Filter by <code>svm.uuid</code>
<code>svm.name</code>	string	query	False	Filter by <code>svm.name</code>
<code>fields</code>	array[string]	query	False	Specify the fields to return.

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

Response

Status: 200, Ok

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

Example response

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

schedule

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

copies

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

Name	Type	Description
retention_period	string	The retention period of Snapshot copies for this schedule. The retention period 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 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.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations

svm

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

snapshot_policy

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

Name	Type	Description
_links	_links	

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Create a Snapshot copy policy

POST /storage/snapshot-policies

Introduced In: 9.6

Creates a Snapshot copy policy.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the Snapshot copy policy.
- `name` - Name for the Snapshot copy policy.
- `copies.schedule` - Schedule at which Snapshot copies are captured on the volume.
- `copies.count` - Number of Snapshot copies to maintain for this schedule.

Recommended optional properties

- `copies.prefix` - Prefix to use when creating Snapshot copies at regular intervals.
- `copies.snapmirror` - Label for SnapMirror operations.
- `copies.retention` - Retention period for Snapshot copy locking enabled volumes.

Default property values

If not specified in POST, the following default property values are assigned:

- `enabled` - *true*
- `copies.prefix` - Value of `schedule.name`

Related ONTAP commands

- `snapshot policy create`

Learn more

- [DOC /storage/snapshot-policies](#)

Parameters

Name	Type	In	Required	Description
<code>return_records</code>	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
<code>_links</code>	_links	
<code>comment</code>	string	A comment associated with the Snapshot copy policy.
<code>copies</code>	array[copies]	

Name	Type	Description
enabled	boolean	Is the Snapshot copy policy enabled?
name	string	Name of the Snapshot copy policy.
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
svm	svm	
uuid	string	

Example request

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

Response

Status: 201, Created

Headers

Name	Description	Type
Location	Useful for tracking the resource location	string

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1638407	When adding schedule to a Snapshot copy policy, the count for that schedule must be specified.
1638408	When adding schedule to a Snapshot copy policy, the schedule name must be specified.
1638413	Schedule not found.
1638417	Specified policy name is invalid.
1638451	This operation would result in total Snapshot copy count for the policy to exceed maximum supported count.
1638508	Another schedule has the same prefix within this policy.
1638526	This operation is not supported on a node Vserver.
1638527	Policy name already exists.
1638528	This operation is not supported in a mixed-version cluster.
1638531	This operation is not supported because specified policy is owned by the cluster admin.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

schedule

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

copies

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

Name	Type	Description
retention_period	string	The retention period of Snapshot copies for this schedule. The retention period 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 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. </num></num></num></num></num>
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations

svm

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

snapshot_policy

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

Name	Type	Description
_links	_links	

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage Snapshot copy policies and schedules

Storage snapshot-policies snapshot_policy.uuid schedules endpoint overview

Overview

In ONTAP, scheduled Snapshot copy creation works based on the schedules associated with Snapshot copy policies. ONTAP provides six cluster-wide schedules: "5min", "8hour", "hourly", "daily", "weekly" and "monthly". A Snapshot copy policy is created using at least one of these schedules and up to 5 schedules can be associated with a Snapshot copy policy. A Snapshot copy policy can be linked to a storage object and based on the schedule in the policy, Snapshot copies are created on the object at that interval. Each schedule in a Snapshot copy policy has a Snapshot copy name prefix attached to it. Every Snapshot copy created using this policy has this prefix in its name. There is also a retention count associated with every schedule. This count indicates the maximum number of Snapshot copies that can exist for a given schedule. Once the Snapshot copy count reaches the retention count, on the next create operation, the oldest Snapshot copy is deleted. A retention period can be associated with every schedule. During Snapshot copy creation, this period is set as SnapLock expiry time on Snapshot copy locking enabled volumes.

A schedule can be added, modified or deleted from a Snapshot copy policy.

Snapshot copy policy schedule APIs

The following APIs are used to perform operations related to Snapshot copy policy schedules:

– POST /api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/

– GET /api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/

– GET /api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

– PATCH /api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

– DELETE /api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

Examples

Adding schedule to a Snapshot copy policy

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

```

# The API:
/api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-818e-11e9-b4f4-005056bbab9c/schedules" -H "accept: application/hal+json" -d '{"schedule.uuid": "7c985d80-818a-11e9-b4f4-005056bbab9c", "count": "5", "prefix": "new_hourly" }'

# The response:
HTTP/1.1 201 Created
Date: Wed, 29 May 2019 22:41:33 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/snapshot-policies/32a0841a-818e-11e9-b4f4-005056bbab9c/schedules
Content-Length: 271
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "snapshot_policy": {
        "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
      },
      "schedule": {
        "uuid": "7c985d80-818a-11e9-b4f4-005056bbab9c"
      },
      "count": 5,
      "prefix": "new_monthly"
    }
  ]
}

```

Retrieving Snapshot copy policy schedules

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

```

# The API:
/api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-818e-11e9-b4f4-005056bbab9c/schedules" -H "accept: application/hal+json"

```

```
# The response:
HTTP/1.1 200 OK
Date: Wed, 29 May 2019 22:49:58 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 898
Content-Type: application/json
{
  "records": [
    {
      "snapshot_policy": {
        "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
      },
      "schedule": {
        "uuid": "63d017dc-818a-11e9-b4f4-005056bbab9c",
        "name": "5min"
      }
    },
    {
      "snapshot_policy": {
        "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
      },
      "schedule": {
        "uuid": "64a5c5da-818a-11e9-b4f4-005056bbab9c",
        "name": "8hour"
      }
    },
    {
      "snapshot_policy": {
        "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
      },
      "schedule": {
        "uuid": "63e21a3e-818a-11e9-b4f4-005056bbab9c",
        "name": "daily"
      }
    },
    {
      "snapshot_policy": {
        "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
      },
      "schedule": {
        "uuid": "7c985d80-818a-11e9-b4f4-005056bbab9c",
        "name": "monthly"
      }
    }
  ]
}
```

```
    }  
  ],  
  "num_records": 4  
}
```

Retrieving the attributes of a specific Snapshot copy policy schedule

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

```
# The API:  
/api/storage/snapshot-  
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}  
  
# The call:  
curl -X GET "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-  
818e-11e9-b4f4-005056bbab9c/schedules/7c985d80-818a-11e9-b4f4-  
005056bbab9c" -H "accept: application/hal+json"  
  
# The response:  
HTTP/1.1 200 OK  
Date: Wed, 29 May 2019 22:54:06 GMT  
Server: libzapid-httpd  
X-Content-Type-Options: nosniff  
Cache-Control: no-cache,no-store,must-revalidate  
Content-Length: 238  
Content-Type: application/json  
{  
  "snapshot_policy": {  
    "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"  
  },  
  "schedule": {  
    "uuid": "7c985d80-818a-11e9-b4f4-005056bbab9c",  
    "name": "monthly"  
  },  
  "count": 5,  
  "prefix": "new_monthly",  
  "snapmirror_label": "-",  
  "retention_period": "PT20M"  
}
```

Updating a Snapshot copy policy schedule

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

```
# The API:
/api/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-
818e-11e9-b4f4-005056bbab9c/schedules/7c985d80-818a-11e9-b4f4-
005056bbab9c" -d '{"count": "10" }' -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 29 May 2019 23:08:00 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
```

Deleting a Snapshot copy policy

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

```
# The API:
/api/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-
818e-11e9-b4f4-005056bbab9c/schedules/7c985d80-818a-11e9-b4f4-
005056bbab9c" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 29 May 2019 23:12:32 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
```

Retrieve Snapshot copy policy schedules

GET /storage/snapshot-policies/{snapshot_policy.uuid}/schedules

Introduced In: 9.8

Retrieves a collection of Snapshot copy policy schedules.

Related ONTAP commands

- `snapshot policy show`

Learn more

- [DOC /storage/snapshot-policies/{snapshot_policy.uuid}/schedules](#)

Parameters

Name	Type	In	Required	Description
snapshot_policy.uuid	string	path	True	Snapshot copy policy UUID
schedule.name	string	query	False	Filter by schedule.name
schedule.uuid	string	query	False	Filter by schedule.uuid
retention_period	string	query	False	Filter by retention_period <ul style="list-style-type: none">• Introduced in: 9.12
count	integer	query	False	Filter by count
prefix	string	query	False	Filter by prefix
snapmirror_label	string	query	False	Filter by snapmirror_label
snapshot_policy.name	string	query	False	Filter by snapshot_policy.name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

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

Response

Status: 200, Ok

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "prefix": "string",
    "schedule": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "weekly",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "snapshot_policy": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "default",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

schedule

Name	Type	Description
_links	_links	
name	string	Job schedule name
uuid	string	Job schedule UUID

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

snapshot_policy_schedule

The Snapshot copy policy schedule object is associated with a Snapshot copy policy and it defines the interval at which Snapshot copies are created and deleted.

Name	Type	Description
_links	_links	

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
retention_period	string	The retention period of Snapshot copies for this schedule.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Add a schedule to a Snapshot copy policy

POST /storage/snapshot-policies/{snapshot_policy.uuid}/schedules

Introduced In: 9.8

Adds a schedule to a Snapshot copy policy.

Required properties

- `schedule.uuid` or `schedule.name` - Schedule at which Snapshot copies are captured on the volume.
- `count` - Number of Snapshot copies to maintain for this schedule.

Recommended optional properties

- `prefix` - Prefix to use when creating Snapshot copies at regular intervals.

Default property values

If not specified in POST, the following default property values are assigned:

- `prefix` - Value of `schedule.name`

Related ONTAP commands

- `snapshot policy add-schedule`

Learn more

- [DOC /storage/snapshot-policies/{snapshot_policy.uuid}/schedules](#)

Parameters

Name	Type	In	Required	Description
<code>snapshot_policy.uuid</code>	string	path	True	Snapshot copy policy UUID
<code>return_records</code>	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
<code>_links</code>	_links	
<code>count</code>	integer	The number of Snapshot copies to maintain for this schedule.
<code>prefix</code>	string	The prefix to use while creating Snapshot copies at regular intervals.

Name	Type	Description
retention_period	string	The retention period of Snapshot copies for this schedule.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "prefix": "string",
  "schedule": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snapshot_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

Response

Status: 201, Created

Headers

Name	Description	Type
Location	Useful for tracking the resource location	string

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1638407	When adding schedule to a Snapshot copy policy, the count for that schedule must be specified.
1638410	Specified schedule already exists in snapshot policy.
1638413	Schedule not found.
1638451	This operation would result in total Snapshot copy count for the policy to exceed maximum supported count.
1638508	Another schedule has the same prefix within this policy.
1638528	This operation is not supported in a mixed-version cluster.
1638531	This operation is not supported because specified policy is owned by the cluster admin.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

schedule

Name	Type	Description
_links	_links	
name	string	Job schedule name
uuid	string	Job schedule UUID

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

snapshot_policy_schedule

The Snapshot copy policy schedule object is associated with a Snapshot copy policy and it defines the interval at which Snapshot copies are created and deleted.

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

Name	Type	Description
retention_period	string	The retention period of Snapshot copies for this schedule.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Delete a schedule from a Snapshot copy policy

```
DELETE /storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}
```

Introduced In: 9.8

Deletes a schedule from a Snapshot copy policy

Related ONTAP commands

- `snapshot policy remove-schedule`

Learn more

- [DOC /storage/snapshot-policies/{snapshot_policy.uuid}/schedules](#)

Parameters

Name	Type	In	Required	Description
snapshot_policy.uuid	string	path	True	Snapshot copy policy UUID
schedule.uuid	string	path	True	Snapshot copy policy schedule UUID

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Code

Error Code	Description
1638412	Schedule does not exist in Snapshot policy.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve Snapshot copy policy schedule details

GET /storage/snapshot-policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

Introduced In: 9.8

Retrieves details of a specific Snapshot copy policy schedule.

Related ONTAP commands

- `snapshot policy show`

Learn more

- [DOC /storage/snapshot-policies/{snapshot_policy.uuid}/schedules](#)

Parameters

Name	Type	In	Required	Description
snapshot_policy.uuid	string	path	True	Snapshot copy policy UUID
schedule.uuid	string	path	True	Snapshot copy policy schedule ID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
retention_period	string	The retention period of Snapshot copies for this schedule.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "prefix": "string",
  "schedule": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snapshot_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

schedule

Name	Type	Description
_links	_links	
name	string	Job schedule name
uuid	string	Job schedule UUID

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

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

Update a Snapshot copy policy schedule

PATCH /storage/snapshot-policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

Introduced In: 9.8

Updates a Snapshot copy policy schedule

Related ONTAP commands

- `snapshot policy modify-schedule`

Learn more

- [DOC /storage/snapshot-policies/{snapshot_policy.uuid}/schedules](#)

Parameters

Name	Type	In	Required	Description
snapshot_policy.uuid	string	path	True	Snapshot copy policy UUID
schedule.uuid	string	path	True	Snapshot copy policy schedule UUID

Request Body

Name	Type	Description
_links	_links	
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
retention_period	string	The retention period of Snapshot copies for this schedule.

Name	Type	Description
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "prefix": "string",
  "schedule": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snapshot_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Code

Error Code	Description
1638451	This operation would result in total Snapshot copy count for the policy to exceed maximum supported count.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

schedule

Name	Type	Description
_links	_links	
name	string	Job schedule name
uuid	string	Job schedule UUID

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

snapshot_policy_schedule

The Snapshot copy policy schedule object is associated with a Snapshot copy policy and it defines the interval at which Snapshot copies are created and deleted.

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

Name	Type	Description
retention_period	string	The retention period of Snapshot copies for this schedule.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Delete a Snapshot copy policy

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

Introduced In: 9.6

Deletes a Snapshot copy policy

Related ONTAP commands

- `snapshot policy delete`

Learn more

- [DOC /storage/snapshot-policies](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Snapshot copy policy UUID

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Code

Error Code	Description
1638415	Cannot delete policy. Reason: Policy is in use by at least one volume.
1638416	Cannot delete policy. Reason: Cannot verify whether policy is in use.
1638430	Cannot delete policy. Reason: Policy is in use by at least one Vserver.
1638430	Cannot delete built-in policy.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve Snapshot copy policy details

GET /storage/snapshot-policies/{uuid}

Introduced In: 9.6

Retrieves details of a specific Snapshot copy policy.

Related ONTAP commands

- `snapshot policy show`

Learn more

- [DOC /storage/snapshot-policies](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Snapshot copy policy UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

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

Example response

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

schedule

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

copies

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

Name	Type	Description
retention_period	string	The retention period of Snapshot copies for this schedule. The retention period 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 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. </num></num></num></num></num>
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations

svm

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Update a Snapshot copy policy

PATCH /storage/snapshot-policies/{uuid}

Introduced In: 9.6

Updates a Snapshot copy policy

Related ONTAP commands

- `snapshot policy modify`
- `snapshot policy modify-schedule`
- `snapshot policy add-schedule`

Learn more

- [DOC /storage/snapshot-policies](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Snapshot copy policy UUID

Request Body

Name	Type	Description
_links	_links	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[copies]	

Name	Type	Description
enabled	boolean	Is the Snapshot copy policy enabled?
name	string	Name of the Snapshot copy policy.
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
svm	svm	
uuid	string	

Example request

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


Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Code

Error Code	Description
1638414	Cannot enable policy. Reason: Specified schedule not found.

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

schedule

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

copies

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

Name	Type	Description
retention_period	string	The retention period of Snapshot copies for this schedule. The retention period 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 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. </num></num></num></num></num>
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations

svm

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

snapshot_policy

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

Name	Type	Description
_links	_links	

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage storage switches

Storage switches endpoint overview

Retrieving storage switch information

The storage switch GET API retrieves all of the switches in the cluster.

Examples

1) Retrieves a list of storage switches from the cluster

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

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

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

# The response:
{
  "records": [
    {
      "name": "Brocade_10.226.57.206",
      "_links": {
        "self": {
          "href": "/api/storage/switches/Brocade_10.226.57.206"
        }
      }
    },
    {
      "name": "Brocade_10.226.57.207",
      "_links": {
        "self": {
          "href": "/api/storage/switches/Brocade_10.226.57.207"
        }
      }
    },
    {
      "name": "Brocade_10.226.57.208",
      "_links": {
        "self": {
          "href": "/api/storage/switches/Brocade_10.226.57.208"
        }
      }
    },
  ],
}
```

```

    {
      "name": "Brocade_10.226.57.209",
      "_links": {
        "self": {
          "href": "/api/storage/switches/Brocade_10.226.57.209"
        }
      }
    }
  ],
  "num_records": 4,
  "_links": {
    "self": {
      "href": "/api/storage/switches/"
    }
  }
}

```

2) Retrieves a specific storage switch from the cluster

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

```

# The API:
/api/storage/switches/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/switches/Brocade_10.226.57.206"
-H "accept: application/hal+json"

# The response:
{
  "name": "Brocade_10.226.57.206",
  "domain_id": 5,
  "switch_fabric_name": "100050eb1a238892",
  "fw_version": "v7.2.1c1",
  "ip_address": "10.226.57.206",
  "is_director_class": false,
  "local": false,
  "monitoring_enabled": true,
  "model": "Brocade6510",
  "role": "subordinate",
  "state": "ok",
  "symbolic_name": "rtp-fc01-41kk11",

```

```
"vendor": "brocade",
"wwn": "100050eb1a1ef7d7",
"power_supply_units": [
  {
    "name": "Power Supply #1",
    "state": "ok"
  },
  {
    "name": "Power Supply #2",
    "state": "ok"
  }
],
"temperature_sensors": [
  {
    "name": "SLOT #0: TEMP #1",
    "reading": 52,
    "state": "ok"
  }
],
"ports": [
  {
    "name": "FC port 0/0",
    "mode": "f_port",
    "wwn": "200050eb1a1ef7d7",
    "enabled": true,
    "state": "online",
    "speed": 16,
    "sfp": {
      "type": "small_form_factor",
      "transmitter_type": "short_wave_laser",
      "serial_number": "HAA2140310058E5"
    }
  },
  {
    "name": "FC port 0/1",
    "mode": "f_port",
    "wwn": "200050eb1a1ef2d7",
    "enabled": true,
    "state": "online",
    "speed": 16,
    "sfp": {
      "type": "small_form_factor",
      "transmitter_type": "short_wave_laser",
      "serial_number": "HAA2140310058E5"
    }
  }
],
```

```
{
  "name": "FC port 0/2",
  "mode": "f_port",
  "wwn": "200050eb1a1ef7d0",
  "enabled": true,
  "state": "online",
  "speed": 16,
  "sfp": {
    "type": "small_form_factor",
    "transmitter_type": "short_wave_laser",
    "serial_number": "HAA2140310058E5"
  }
},
{
  "name": "FC port 0/3",
  "mode": "f_port",
  "wwn": "200050eb1a1ef7d7",
  "enabled": true,
  "state": "online",
  "speed": 16,
  "sfp": {
    "type": "small_form_factor",
    "transmitter_type": "short_wave_laser",
    "serial_number": "HAA2140310058E5"
  }
},
{
  "name": "FC port 0/4",
  "mode": "f_port",
  "wwn": "200050eb1a1ef2d7",
  "enabled": true,
  "state": "online",
  "speed": 16,
  "sfp": {
    "type": "small_form_factor",
    "transmitter_type": "short_wave_laser",
    "serial_number": "HAA2140310058E5"
  }
},
{
  "name": "FC port 0/5",
  "mode": "f_port",
  "wwn": "200050eb1a1ef7d0",
  "enabled": true,
  "state": "online",
  "speed": 16,
```



```

    "sfp": {
      "type": "small_form_factor",
      "transmitter_type": "short_wave_laser",
      "serial_number": "HAA2140310058E5"
    }
  ],
  "connections": [
    {
      "source_port": {
        "name": "FC port 0/0",
        "wwn": "200050eb1a236efd",
        "mode": "f_port"
      },
      "peer_port": {
        "wwn": "2100000e1e30ac5f",
        "connection": "sti8020mcc-htp-006:fcvi_device_1",
        "type": "fcvi_adapter",
        "unique_id": "38993dc0-4ea1-11eb-9331-00a0985bd455"
      }
    },
    {
      "source_port": {
        "name": "FC port 0/1",
        "wwn": "200150eb1a236efd",
        "mode": "f_port"
      },
      "peer_port": {
        "wwn": "21000024ff72c0c9",
        "connection": "sti8020mcc-htp-006:2b",
        "type": "fcp_adapter",
        "unique_id": "38993dc0-4ea1-11eb-9331-00a0985bd455"
      }
    },
    {
      "source_port": {
        "name": "FC port 0/2",
        "wwn": "200250eb1a236efd",
        "mode": "f_port"
      },
      "peer_port": {
        "wwn": "21000024ff72c0cb",
        "connection": "sti8020mcc-htp-006:2d",
        "type": "fcp_adapter",
        "unique_id": "38993dc0-4ea1-11eb-9331-00a0985bd455"
      }
    }
  ]
}

```

```

    }
  ],
  "fans": [
    {
      "name": "FAN #1",
      "speed": 7336,
      "state": "ok"
    },
    {
      "name": "FAN #2",
      "speed": 7336,
      "state": "ok"
    }
  ],
  "paths": [
    {
      "adapter": {
        "name": "2a",
        "wwn": "21000024ff6c4bc0",
        "type": "fcp_initiator"
      },
      "port": {
        "name": "FC port 0/4",
        "speed": 8
      },
      "node": {
        "name": "sti8020mcc-htp-005",
        "uuid": "382cb083-4416-11eb-ad1d-00a0985bd455",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/382cb083-4416-11eb-ad1d-00a0985bd455"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/storage/ports/382cb083-4416-11eb-ad1d-00a0985bd455/2a"
        }
      }
    },
    {
      "adapter": {
        "name": "2c",
        "wwn": "21000024ff6c4bc2",

```

```

    "type": "fcv_initiator"
  },
  "port": {
    "name": "FC port 0/5",
    "speed": 8
  },
  "node": {
    "name": "sti8020mcc-htp-005",
    "uuid": "382cb083-4416-11eb-ad1d-00a0985bd455",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/382cb083-4416-11eb-ad1d-
00a0985bd455"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/ports/382cb083-4416-11eb-ad1d-
00a0985bd455/2c"
    }
  }
},
{
  "adapter": {
    "name": "fcvi_device_0",
    "wwn": "2100000e1e09d5d2",
    "type": "fc_vi"
  },
  "port": {
    "name": "FC port 0/3",
    "speed": 16
  },
  "node": {
    "name": "sti8020mcc-htp-005",
    "uuid": "382cb083-4416-11eb-ad1d-00a0985bd455",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/382cb083-4416-11eb-ad1d-
00a0985bd455"
      }
    }
  }
},
{
  "adapter": {

```

```

    "name": "2a",
    "wwn": "21000024ff72c0c8",
    "type": "fc_initiator"
  },
  "port": {
    "name": "FC port 0/1",
    "speed": 8
  },
  "node": {
    "name": "sti8020mcc-htp-006",
    "uuid": "364fbba8-4416-11eb-8e72-00a098431045",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/364fbba8-4416-11eb-8e72-00a098431045"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/ports/364fbba8-4416-11eb-8e72-00a098431045/2a"
    }
  }
},
{
  "adapter": {
    "name": "2c",
    "wwn": "21000024ff72c0ca",
    "type": "fc_initiator"
  },
  "port": {
    "name": "FC port 0/2",
    "speed": 8
  },
  "node": {
    "name": "sti8020mcc-htp-006",
    "uuid": "364fbba8-4416-11eb-8e72-00a098431045",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/364fbba8-4416-11eb-8e72-00a098431045"
      }
    }
  },
  "_links": {

```

```

      "self": {
        "href": "/api/storage/ports/364fbba8-4416-11eb-8e72-
00a098431045/2c"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/switches/Brocade_10.226.57.206"
    }
  }
}

```

Retrieve storage switches

GET /storage/switches

Introduced In: 9.9

Retrieves a collection of storage switches.

Related ONTAP commands

- `storage switch show`

Learn more

- [DOC /storage/switches](#)

Parameters

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • 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[storage_switch]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "connections": {
      "peer_port": {
        "type": "unknown"
      }
    },
    "errors": {
      "reason": {
        "arguments": {
          "code": "string",
          "message": "string"
        },
        "code": "4",
        "message": "entry doesn't exist",
        "target": "uuid"
      },
      "severity": "unknown",
      "type": "switch_unreachable"
    },
    "fans": {
      "state": "ok"
    },
    "monitored_blades": {
    },
    "paths": {
      "adapter": {
        "type": "unknown"
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      }
    },
  },
}
```

```

    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"ports": {
  "mode": "unknown",
  "sfp": {
    "transmitter_type": "unknown",
    "type": "unknown"
  },
  "state": "error"
},
"power_supply_units": {
  "state": "ok"
},
"role": "unknown",
"state": "ok",
"temperature_sensors": {
  "state": "error"
},
"vendor": "unknown",
"vsans": {
  "state": "ok"
},
"zones": {
}
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

peer_port

Name	Type	Description
connection	string	Storage switch peer port host and name
type	string	Storage switch peer type
unique_id	string	Storage switch peer unique ID
wwn	string	Storage switch peer port world wide name

source_port

Name	Type	Description
mode	string	Storage switch port operating mode
name	string	Storage switch port name
wwn	string	Storage switch peer port world wide name

connections

Name	Type	Description
peer_port	peer_port	
source_port	source_port	

component

Name	Type	Description
id	integer	Error component ID
name	string	Error component name

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

errors

Name	Type	Description
component	component	
reason	error	
severity	string	Error component severity
type	string	Error component type

fans

Name	Type	Description
name	string	Storage switch fan name
speed	integer	Storage switch fan speed
state	string	Storage switch fan state

adapter

Name	Type	Description
name	string	Node adapter name
type	string	Node adapter type
wwn	string	Node adapter world wide name

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

port

Name	Type	Description
name	string	Storage switch port name
speed	integer	Storage switch port speed, in Gbps

paths

Name	Type	Description
adapter	adapter	
node	node	
port	port	

sfp

Name	Type	Description
serial_number	string	Storage switch port SFP serial number

Name	Type	Description
transmitter_type	string	Storage switch port SFP transmitter type
type	string	Storage switch port SFP type

ports

Name	Type	Description
enabled	boolean	Indicates whether the storage switch port is enabled.
mode	string	Storage switch port mode
name	string	Storage switch port name
sfp	sfp	
speed	integer	Storage switch port speed, in Gbps
state	string	Storage switch port state
wwn	string	Storage switch port world wide name

power_supply_units

Name	Type	Description
name	string	Power supply unit name
state	string	Power supply unit state

temperature_sensors

Name	Type	Description
name	string	Temperature sensor name
reading	integer	Temperature sensor reading, in degrees celsius.
state	string	Temperature sensor state

vsans

Name	Type	Description
id	integer	Storage switch VSAN ID
iod	boolean	Indicates whether in-order delivery is set for a zone.
load_balancing_types	string	Storage switch VSAN load balancing type
name	string	Storage switch VSAN name
state	string	Storage switch VSAN Port state

port

Name	Type	Description
id	string	Storage switch zone port ID
name	string	Storage switch zone port

zones

Name	Type	Description
id	integer	Storage switch zone ID
name	string	Storage switch zone name
port	port	
wwn	string	Storage switch zone world wide name

storage_switch

The Storage switch object describes the storage switch properties, features and cabling.

Name	Type	Description
connections	array[connections]	
director_class	boolean	
domain_id	integer	Domain ID
errors	array[errors]	

Name	Type	Description
fabric_name	string	Storage switch fabric name
fans	array[fans]	
firmware_version	string	Storage switch firmware version
ip_address	string	IP Address
local	boolean	Indicates whether the storage switch is directly connected to the reporting cluster.
model	string	Storage switch model.
monitored_blades	array[integer]	Indicates the blades that are being monitored for a director-class switch.
monitoring_enabled	boolean	Indicates whether monitoring is enabled for the storage switch.
name	string	Storage switch name
paths	array[paths]	
ports	array[ports]	
power_supply_units	array[power_supply_units]	
role	string	Storage switch role in fabric.
state	string	Storage switch state
symbolic_name	string	Storage switch symbolic name
temperature_sensors	array[temperature_sensors]	
vendor	string	Storage switch vendor
vsans	array[vsans]	
wwn	string	Storage switch world wide name
zones	array[zones]	

Retrieve a specific storage switch

GET /storage/switches/{name}

Introduced In: 9.9

Retrieves a specific storage switch.

Related ONTAP commands

- `storage switch show`

Learn more

- [DOC /storage/switches](#)

Parameters

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

Response

```
Status: 200, Ok
```

Name	Type	Description
connections	array[connections]	
director_class	boolean	
domain_id	integer	Domain ID
errors	array[errors]	
fabric_name	string	Storage switch fabric name
fans	array[fans]	
firmware_version	string	Storage switch firmware version
ip_address	string	IP Address
local	boolean	Indicates whether the storage switch is directly connected to the reporting cluster.
model	string	Storage switch model.

Name	Type	Description
monitored_blades	array[integer]	Indicates the blades that are being monitored for a director-class switch.
monitoring_enabled	boolean	Indicates whether monitoring is enabled for the storage switch.
name	string	Storage switch name
paths	array[paths]	
ports	array[ports]	
power_supply_units	array[power_supply_units]	
role	string	Storage switch role in fabric.
state	string	Storage switch state
symbolic_name	string	Storage switch symbolic name
temperature_sensors	array[temperature_sensors]	
vendor	string	Storage switch vendor
vsans	array[vsans]	
wwn	string	Storage switch world wide name
zones	array[zones]	

Example response

```
{
  "connections": {
    "peer_port": {
      "type": "unknown"
    }
  },
  "errors": {
    "reason": {
      "arguments": {
        "code": "string",
        "message": "string"
      },
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    },
    "severity": "unknown",
    "type": "switch_unreachable"
  },
  "fans": {
    "state": "ok"
  },
  "monitored_blades": {
  },
  "paths": {
    "adapter": {
      "type": "unknown"
    },
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "ports": {
    "mode": "unknown",
    "sfp": {
      "transmitter_type": "unknown",
      "type": "unknown"
    }
  },
}
```

```

    "state": "error"
  },
  "power_supply_units": {
    "state": "ok"
  },
  "role": "unknown",
  "state": "ok",
  "temperature_sensors": {
    "state": "error"
  },
  "vendor": "unknown",
  "vsans": {
    "state": "ok"
  },
  "zones": {
  }
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

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

```

Definitions

See Definitions

peer_port

Name	Type	Description
connection	string	Storage switch peer port host and name
type	string	Storage switch peer type
unique_id	string	Storage switch peer unique ID
wwn	string	Storage switch peer port world wide name

source_port

Name	Type	Description
mode	string	Storage switch port operating mode
name	string	Storage switch port name
wwn	string	Storage switch peer port world wide name

connections

Name	Type	Description
peer_port	peer_port	
source_port	source_port	

component

Name	Type	Description
id	integer	Error component ID
name	string	Error component name

error_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

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

errors

Name	Type	Description
component	component	
reason	error	
severity	string	Error component severity
type	string	Error component type

fans

Name	Type	Description
name	string	Storage switch fan name
speed	integer	Storage switch fan speed
state	string	Storage switch fan state

adapter

Name	Type	Description
name	string	Node adapter name
type	string	Node adapter type
wwn	string	Node adapter world wide name

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

port

Name	Type	Description
name	string	Storage switch port name
speed	integer	Storage switch port speed, in Gbps

paths

Name	Type	Description
adapter	adapter	
node	node	
port	port	

sfp

Name	Type	Description
serial_number	string	Storage switch port SFP serial number
transmitter_type	string	Storage switch port SFP transmitter type
type	string	Storage switch port SFP type

ports

Name	Type	Description
enabled	boolean	Indicates whether the storage switch port is enabled.
mode	string	Storage switch port mode
name	string	Storage switch port name
sfp	sfp	
speed	integer	Storage switch port speed, in Gbps
state	string	Storage switch port state
wwn	string	Storage switch port world wide name

power_supply_units

Name	Type	Description
name	string	Power supply unit name
state	string	Power supply unit state

temperature_sensors

Name	Type	Description
name	string	Temperature sensor name
reading	integer	Temperature sensor reading, in degrees celsius.
state	string	Temperature sensor state

vsans

Name	Type	Description
id	integer	Storage switch VSAN ID
iod	boolean	Indicates whether in-order delivery is set for a zone.

Name	Type	Description
load_balancing_types	string	Storage switch VSAN load balancing type
name	string	Storage switch VSAN name
state	string	Storage switch VSAN Port state

port

Name	Type	Description
id	string	Storage switch zone port ID
name	string	Storage switch zone port

zones

Name	Type	Description
id	integer	Storage switch zone ID
name	string	Storage switch zone name
port	port	
wwn	string	Storage switch zone world wide name

Manage storage tape devices

Storage tape-devices endpoint overview

Retrieving storage tape information

The storage tape GET API retrieves all of the tapes in the cluster.

Examples

1) Retrieving a list of tapes from the cluster

The following example returns the list of tapes in the cluster:

```

# The API:
/api/storage/tape-devices

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

# The response:
{
  "records": [
    {
      "node": {
        "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
        "name": "st-8020-1-01",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-
00a0985ebbe7"
          }
        }
      },
      "device_id": "2d.0",
      "_links": {
        "self": {
          "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-
00a0985ebbe7/2d.0"
        }
      }
    },
    {
      "node": {
        "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
        "name": "st-8020-1-01",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-
00a0985ebbe7"
          }
        }
      },
      "device_id": "2d.0L1",
      "_links": {
        "self": {
          "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-
00a0985ebbe7/2d.0L1"
        }
      }
    }
  ]
}

```

```

    }
  },
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-00a0985ebbe7"
        }
      }
    },
    "device_id": "qeg-tape-brocade2-8g:0.126",
    "_links": {
      "self": {
        "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-00a0985ebbe7/qeg-tape-brocade2-8g%3A0.126"
      }
    }
  },
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-00a0985ebbe7"
        }
      }
    },
    "device_id": "stsw-broc6510-01:11.126",
    "_links": {
      "self": {
        "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-00a0985ebbe7/stsw-broc6510-01%3A11.126"
      }
    }
  },
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {

```

```

        "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-00a0985ebbe7"
      }
    },
    "device_id": "stsw-broc6510-01:15.126",
    "_links": {
      "self": {
        "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-00a0985ebbe7/stsw-broc6510-01%3A15.126"
      }
    }
  },
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-00a0985ebbe7"
        }
      }
    },
    "device_id": "stsw-broc6510-01:15.126L1",
    "_links": {
      "self": {
        "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-00a0985ebbe7/stsw-broc6510-01%3A15.126L1"
      }
    }
  },
  {
    "node": {
      "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
      "name": "st-8020-1-01",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-00a0985ebbe7"
        }
      }
    },
    "device_id": "stsw-broc6510-01:22.126",
    "_links": {
      "self": {

```

```

    "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-00a0985ebbe7/stsw-broc6510-01%3A22.126"
  }
}
},
{
  "node": {
    "uuid": "4083be52-5315-11eb-a839-00a0985ebbe7",
    "name": "st-8020-1-01",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/4083be52-5315-11eb-a839-00a0985ebbe7"
      }
    }
  },
  "device_id": "stsw-broc6510-01:23.126",
  "_links": {
    "self": {
      "href": "/api/storage/tape-devices/4083be52-5315-11eb-a839-00a0985ebbe7/stsw-broc6510-01%3A23.126"
    }
  }
},
],
"num_records": 7,
"_links": {
  "self": {
    "href": "/api/storage/tape-devices"
  }
}
}
}

```

2) Retrieving a specific tape device from the cluster

The following example returns the requested tape device. If there is no tape with the requested UID, an error is returned.

```

# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-

```

```
11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json"

# The response:
{
  "node": {
    "uuid": "5f5275eb-5315-11eb-8ac4-00a0985e0dcf",
    "name": "st-8020-1-02",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/5f5275eb-5315-11eb-8ac4-00a0985e0dcf"
      }
    }
  },
  "device_id": "2d.0",
  "type": "tape",
  "description": "IBM LTO-6 ULT3580",
  "alias": {
    "name": "st7",
    "mapping": "SN[1068000245]"
  },
  "aliases": [
    {
      "name": "st7",
      "mapping": "SN[1068000245]"
    }
  ],
  "wwnn": "5001697722ee0010",
  "wwpn": "5001697722ee0011",
  "serial_number": "1068000245",
  "interface": "sas",
  "device_state": "offline",
  "formats": [
    "LTO-4/5 Native Density",
    "LTO-4/5 Compressed",
    "LTO-6 2.5TB",
    "LTO-6 6.25TB Compressed"
  ],
  "storage_port": {
    "name": "2d"
  },
  "file_number": -1,
  "block_number": -1,
  "residual_count": 0,
  "density": "low",
  "device_names": [
    {
```

```
"rewind_device": "rst0l",
"no_rewind_device": "nrst0l",
"unload_reload_device": "urst0l"
},
{
"rewind_device": "rst0m",
"no_rewind_device": "nrst0m",
"unload_reload_device": "urst0m"
},
{
"rewind_device": "rst0h",
"no_rewind_device": "nrst0h",
"unload_reload_device": "urst0h"
},
{
"rewind_device": "rst0a",
"no_rewind_device": "nrst0a",
"unload_reload_device": "urst0a"
}
],
"reservation_type": "off",
"_links": {
"self": {
"href": "/api/storage/tape-devices/5f5275eb-5315-11eb-8ac4-00a0985e0dcf/2d.0"
}
}
}
```

Updating a tape device

The tape PATCH API allows the tape device to be set online or offline, positioned, and given an alias.

Examples

1) Taking a tape device offline

The following example takes a tape device offline:

```
# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json" -H "Content-Type: application/hal_json" -d '{"online": "false"}'

# The response:
{
  "job": {
    "uuid": "9e544626-306e-11ec-8c2e-00a098b81daa",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/9e544626-306e-11ec-8c2e-00a098b81daa"
      }
    }
  }
}
```

2) Bringing a tape device online

The following example brings a tape device online:

```
# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json" -H "Content-Type: application/hal_json" -d '{"online": "true"}'

# The response:
{
  "job": {
    "uuid": "9e544626-306e-11ec-8c2e-00a098b81daa",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/9e544626-306e-11ec-8c2e-00a098b81daa"
      }
    }
  }
}
```

3) Giving a tape device an alias

The following example assigns an alias to a tape device:

```
# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json" -H "Content-Type: application/hal_json" -d '{"aliases": [{"name": "st0"}]}'

# The response:
{
}
```

4) Removing a tape device's aliases

The following example clears any aliases previously assigned to a tape device:

```
# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json" -H "Content-Type: application/hal_json" -d '{"aliases": []}'

# The response:
{
}
```

5) Rewinding a tape device

The following example rewinds a tape device:

```
# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json" -H "Content-Type: application/hal_json" -d '{"position": {"operation": "rewind"}}'

# The response:
{
  "job": {
    "uuid": "9e544626-306e-11ec-8c2e-00a098b81daa",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/9e544626-306e-11ec-8c2e-00a098b81daa"
      }
    }
  }
}
```

6) Forwarding the tape five files

The following example moves the tape forward five file records:

```

# The API:
/api/storage/tape-devices/{node.uuid}/{device_id}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/tape-devices/5f5275eb-5315-11eb-8ac4-00a0985e0dcf/2d.0" -H "accept: application/hal+json" -H "Content-Type: application/hal_json" -d '{"position": {"operation": "fsf", "count": 5}}'

# The response:
{
  "job": {
    "uuid": "954c20d5-306e-11ec-8c2e-00a098b81daa",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/954c20d5-306e-11ec-8c2e-00a098b81daa"
      }
    }
  }
}

```

Retrieve tape devices

GET /storage/tape-devices

Introduced In: 9.9

Retrieves a collection of tape devices.

Related ONTAP commands

- `storage tape show`

Learn more

- [DOC /storage/tape-devices](#)

Parameters

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

Name	Type	In	Required	Description
wwpn	string	query	False	Filter by wwpn
device_id	string	query	False	Filter by device_id
type	string	query	False	Filter by type
device_names.no_rewind_device	string	query	False	Filter by device_names.no_rewind_device
device_names.unload_reload_device	string	query	False	Filter by device_names.unload_reload_device
device_names.rewind_device	string	query	False	Filter by device_names.rewind_device
residual_count	integer	query	False	Filter by residual_count
wwnn	string	query	False	Filter by wwnn
aliases.name	string	query	False	Filter by aliases.name • Introduced in: 9.11
aliases.mapping	string	query	False	Filter by aliases.mapping • Introduced in: 9.11
alias.mapping	string	query	False	Filter by alias.mapping
alias.name	string	query	False	Filter by alias.name
file_number	integer	query	False	Filter by file_number
block_number	integer	query	False	Filter by block_number
interface	string	query	False	Filter by interface

Name	Type	In	Required	Description
serial_number	string	query	False	Filter by serial_number
description	string	query	False	Filter by description
online	boolean	query	False	Filter by online <ul style="list-style-type: none"> • Introduced in: 9.11
formats	string	query	False	Filter by formats
reservation_type	string	query	False	Filter by reservation_type
device_state	string	query	False	Filter by device_state
density	string	query	False	Filter by density <ul style="list-style-type: none"> • Introduced in: 9.11
storage_port.name	string	query	False	Filter by storage_port.name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • 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[tape_device]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "alias": {
      "mapping": "SN[10WT000933]",
      "name": "st6"
    },
    "aliases": {
      "mapping": "SN[10WT000933]",
      "name": "st6"
    },
    "block_number": 0,
    "density": "low",
    "description": "QUANTUM LTO-8 ULTRIUM",
    "device_id": "1a.0",
    "device_names": {
      "no_rewind_device": "nrst6l",
      "rewind_device": "rst6l",
      "unload_reload_device": "urst6l"
    },
    "device_state": "read_write_enabled",
    "file_number": 0,
    "formats": [
      "LTO-7 6TB",
      "LTO-7 15TB Compressed",
      "LTO-8 12TB",
      "LTO-8 30TB Compressed"
    ],
    "interface": "sas",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",

```

```

    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "position": {
    "count": 5,
    "operation": "rewind"
  },
  "reservation_type": "off",
  "residual_count": 0,
  "serial_number": "10WT00093",
  "storage_port": {
    "name": "2b"
  },
  "type": "tape",
  "wwnn": "500507631295741c",
  "wwpn": "500507631295741c"
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

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

```


Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

alias

Name	Type	Description
mapping	string	This field will no longer be supported in a future release. Use <code>aliases.mapping</code> instead.
name	string	This field will no longer be supported in a future release. Use <code>aliases.name</code> instead.

aliases

Name	Type	Description
mapping	string	Alias mapping.
name	string	Alias name.

device_names

Name	Type	Description
no_rewind_device	string	Device name for no rewind.
rewind_device	string	Device name for rewind.
unload_reload_device	string	Device name for unload or reload operations.

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

position

Name	Type	Description
count	integer	Number of times to run position operation.
operation	string	Position operation.

storage_port

Name	Type	Description
name	string	Initiator port.

tape_device

Name	Type	Description
alias	alias	
aliases	array[aliases]	
block_number	integer	Block number.
density	string	Density.
description	string	
device_id	string	
device_names	array[device_names]	
device_state	string	Operational state of the device.
file_number	integer	File number.
formats	array[string]	Tape cartridge format.

Name	Type	Description
interface	string	Device interface type.
node	node	
online	boolean	
position	position	
reservation_type	string	
residual_count	integer	Residual count of the last I/O operation.
serial_number	string	
storage_port	storage_port	
type	string	Device type.
wwnn	string	World Wide Node Name.
wwpn	string	World Wide Port Name.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a tape device

GET /storage/tape-devices/{node.uuid}/{device_id}

Introduced In: 9.9

Retrieves a specific tape device.

Related ONTAP commands

- `storage tape show`

Learn more

- [DOC /storage/tape-devices](#)

Parameters

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

Response

Status: 200, Ok

Name	Type	Description
alias	alias	
aliases	array[aliases]	
block_number	integer	Block number.
density	string	Density.
description	string	
device_id	string	
device_names	array[device_names]	
device_state	string	Operational state of the device.
file_number	integer	File number.
formats	array[string]	Tape cartridge format.
interface	string	Device interface type.

Name	Type	Description
node	node	
online	boolean	
position	position	
reservation_type	string	
residual_count	integer	Residual count of the last I/O operation.
serial_number	string	
storage_port	storage_port	
type	string	Device type.
wwnn	string	World Wide Node Name.
wwpn	string	World Wide Port Name.

Example response

```
{
  "alias": {
    "mapping": "SN[10WT000933]",
    "name": "st6"
  },
  "aliases": {
    "mapping": "SN[10WT000933]",
    "name": "st6"
  },
  "block_number": 0,
  "density": "low",
  "description": "QUANTUM LTO-8 ULTRIUM",
  "device_id": "1a.0",
  "device_names": {
    "no_rewind_device": "nrst6l",
    "rewind_device": "rst6l",
    "unload_reload_device": "urst6l"
  },
  "device_state": "read_write_enabled",
  "file_number": 0,
  "formats": [
    "LTO-7 6TB",
    "LTO-7 15TB Compressed",
    "LTO-8 12TB",
    "LTO-8 30TB Compressed"
  ],
  "interface": "sas",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "position": {
    "count": 5,
    "operation": "rewind"
  },
  "reservation_type": "off",
  "residual_count": 0,
  "serial_number": "10WT00093",
  "storage_port": {
```

```
  "name": "2b"
},
"type": "tape",
"wwnn": "500507631295741c",
"wwpn": "500507631295741c"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

alias

Name	Type	Description
mapping	string	This field will no longer be supported in a future release. Use aliases.mapping instead.
name	string	This field will no longer be supported in a future release. Use aliases.name instead.

aliases

Name	Type	Description
mapping	string	Alias mapping.
name	string	Alias name.

device_names

Name	Type	Description
no_rewind_device	string	Device name for no rewind.
rewind_device	string	Device name for rewind.
unload_reload_device	string	Device name for unload or reload operations.

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	
uuid	string	

position

Name	Type	Description
count	integer	Number of times to run position operation.
operation	string	Position operation.

storage_port

Name	Type	Description
name	string	Initiator port.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Update a tape device

PATCH /storage/tape-devices/{node.uuid}/{device_id}

Introduced In: 9.11

Updates a specific tape device.

Related ONTAP commands

- `storage tape alias-set`
- `storage tape alias-clear`
- `storage tape online`
- `storage tape offline`
- `storage tape position`

Learn more

- [DOC /storage/tape-devices](#)

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Node UUID
device_id	string	path	True	Device ID

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

Request Body

Name	Type	Description
alias	alias	
aliases	array[aliases]	
block_number	integer	Block number.
density	string	Density.
description	string	
device_id	string	
device_names	array[device_names]	

Name	Type	Description
device_state	string	Operational state of the device.
file_number	integer	File number.
formats	array[string]	Tape cartridge format.
interface	string	Device interface type.
node	node	
online	boolean	
position	position	
reservation_type	string	
residual_count	integer	Residual count of the last I/O operation.
serial_number	string	
storage_port	storage_port	
type	string	Device type.
wwnn	string	World Wide Node Name.
wwpn	string	World Wide Port Name.

Example request

```
{
  "alias": {
    "mapping": "SN[10WT000933]",
    "name": "st6"
  },
  "aliases": {
    "mapping": "SN[10WT000933]",
    "name": "st6"
  },
  "block_number": 0,
  "density": "low",
  "description": "QUANTUM LTO-8 ULTRIUM",
  "device_id": "1a.0",
  "device_names": {
    "no_rewind_device": "nrst6l",
    "rewind_device": "rst6l",
    "unload_reload_device": "urst6l"
  },
  "device_state": "read_write_enabled",
  "file_number": 0,
  "formats": [
    "LTO-7 6TB",
    "LTO-7 15TB Compressed",
    "LTO-8 12TB",
    "LTO-8 30TB Compressed"
  ],
  "interface": "sas",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "position": {
    "count": 5,
    "operation": "rewind"
  },
  "reservation_type": "off",
  "residual_count": 0,
  "serial_number": "10WT00093",
  "storage_port": {
```

```

    "name": "2b"
  },
  "type": "tape",
  "wwnn": "500507631295741c",
  "wwpn": "500507631295741c"
}

```

Response

Status: 200, Ok

Response

Status: 202, Accepted

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
11403264	Tape operation "<operation>" failed on tape device \"<device_id>\". This may indicate tape reached end of data, a hardware error, an illegal request, an invalid name format, or an aborted command.</device_id></operation>
11403265	Tape operation "<operation>" failed because it encountered an unexpected file mark on tape device \"<device_id>\".</device_id></operation>
11403266	Internal error. Unknown tape parameter for tape device "<device_id>\".</device_id>
11403267	Tape device "<device_id>" is busy.</device_id>
11403268	No tape loaded for "<operation>" operation on tape drive \"<device_id>\".</device_id></operation>
11403269	Tape device "<device_id>" unknown.</device_id>
11403270	Tape alias name "<alias>" already exists.</alias>
11403271	The format of the tape alias name "<alias>" is invalid. Use \"st\" or \"mcl\" followed by one or more digits.</alias>

Error Code	Description
11403273	Internal error. Tape alias name "<alias>" not saved.</alias>
11403274	Internal error. Tape alias operation "<operation>" for alias name "<alias>" failed.</alias></operation>
11403275	Tape alias name "<alias>" not found.</alias>
11403277	Tape device "<device_id>" is reserved by another host.</device_id>
11403278	Invalid count value specified with operation "<operation>". Valid values are from 0 to {1}.</operation>

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

alias

Name	Type	Description
mapping	string	This field will no longer be supported in a future release. Use aliases.mapping instead.
name	string	This field will no longer be supported in a future release. Use aliases.name instead.

aliases

Name	Type	Description
mapping	string	Alias mapping.
name	string	Alias name.

device_names

Name	Type	Description
no_rewind_device	string	Device name for no rewind.
rewind_device	string	Device name for rewind.
unload_reload_device	string	Device name for unload or reload operations.

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	
uuid	string	

position

Name	Type	Description
count	integer	Number of times to run position operation.
operation	string	Position operation.

storage_port

Name	Type	Description
name	string	Initiator port.

tape_device

Name	Type	Description
alias	alias	
aliases	array[aliases]	
block_number	integer	Block number.
density	string	Density.
description	string	
device_id	string	
device_names	array[device_names]	
device_state	string	Operational state of the device.
file_number	integer	File number.
formats	array[string]	Tape cartridge format.
interface	string	Device interface type.
node	node	
online	boolean	
position	position	

Name	Type	Description
reservation_type	string	
residual_count	integer	Residual count of the last I/O operation.
serial_number	string	
storage_port	storage_port	
type	string	Device type.
wwnn	string	World Wide Node Name.
wwpn	string	World Wide Port Name.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage volume efficiency policies

Storage volume-efficiency-policies endpoint overview

Overview

Volume efficiency policies specify information about efficiency policies that are applied to the volume.

Volume efficiency policy APIs

The following APIs are used to perform operations related to volume efficiency policy information:

– POST /api/storage/volume-efficiency-policies

– GET /api/storage/volume-efficiency-policies

– GET /api/storage/volume-efficiency-policies/{uuid}

– PATCH /api/storage/volume-efficiency-policies/{uuid}

– DELETE /api/storage/volume-efficiency-policies/{uuid}

Examples

Creating a volume efficiency policy

The POST operation is used to create a volume efficiency policy with the specified attributes.

```

# The API:
/api/storage/volume-efficiency-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volume-efficiency-policies"
-H "accept: application/hal+json" -d '{"name": "new_policy", "type":
"scheduled", "schedule": { "name": "daily" }, "duration": "2",
"qos_policy": "best_effort", "enabled": "true", "comment": "schedule-
policy", "svm": { "name": "vs1"}}'

# The response:
HTTP/1.1 201 Created
Date: Tue, 12 Mar 2019 21:20:24 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/volume-efficiency-policies/a69d8173-450c-11e9-aa44-
005056bbc848
Content-Length: 369
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "uuid": "a69d8173-450c-11e9-aa44-005056bbc848",
      "svm": {
        "name": "vs1"
      },
      "name": "new_policy",
      "type": "scheduled",
      "schedule":{
        "name": "daily"
      },
      "duration": "2",
      "qos_policy": "best_effort",
      "enabled": "true",
      "comment": "schedule-policy"
    }
  ]
}

```

Retrieving volume efficiency policy attributes

The GET operation is used to retrieve volume efficiency policy attributes.

```
# The API:
/api/storage/volume-efficiency-policies

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

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:17:17 GMT
Server: libzapid-http
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 686
Content-Type: application/json
{
"records": [
  {
    "uuid": "3c112527-2fe8-11e9-b55e-005056bbf1c8",
    "name": "default",
    "_links": {
      "self": {
        "href": "/api/storage/volume-efficiency-policies/3c112527-2fe8-
11e9-b55e-005056bbf1c8"
      }
    }
  },
  {
    "uuid": "3c1c1656-2fe8-11e9-b55e-005056bbf1c8",
    "name": "default-1weekly",
    "_links": {
      "self": {
        "href": "/api/storage/volume-efficiency-policies/3c1c1656-2fe8-
11e9-b55e-005056bbf1c8"
      }
    }
  },
  {
    "uuid": "3c228b82-2fe8-11e9-b55e-005056bbf1c8",
    "name": "none",
    "_links": {
      "self": {
        "href": "/api/storage/volume-efficiency-policies/3c228b82-2fe8-
11e9-b55e-005056bbf1c8"
      }
    }
  }
]
```

```

    }
  ],
  "num_records": 3,
  "_links": {
    "self": {
      "href": "/api/storage/volume-efficiency-policies/"
    }
  }
}
}

```

Retrieving the attributes of a specific volume efficiency policy

The GET operation is used to retrieve the attributes of a specific volume efficiency policy.

The API:

/api/storage/volume-efficiency-policies/{uuid}

The call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volume-efficiency-policies/3c112527-2fe8-11e9-b55e-005056bbf1c8" -H "accept: application/hal+json"</mgmt-ip>
```

The response:

```
HTTP/1.1 200 OK Date: Tue, 12 Mar 2019 21:24:48 GMT Server: libzapid-httpd X-Content-Type-Options: nosniff Cache-Control: no-cache,no-store,must-revalidate Content-Length: 381 Content-Type: application/json {
  "uuid": "3c112527-2fe8-11e9-b55e-005056bbf1c8", "name": "new_policy", "type": "scheduled", "schedule": {
    "name": "daily" } "duration": "2", "qos_policy": "best_effort", "enabled": "true", "comment": "schedule-policy",
  "svm": { "name": "vs1" } "_links": { "self": { "href": "/api/storage/volume-efficiency-policies/3c112527-2fe8-11e9-b55e-005056bbf1c8" } } }
```

```
### Updating a volume efficiency policy
```

The PATCH operation is used to update the specific attributes of a volume efficiency policy.

The API:

/api/storage/volume-efficiency-policies/{uuid}

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/storage/volume-efficiency-policies/ae9e65c4-4506-11e9-aa44-005056bbc848" -d '{"duration": "3"}' -H "accept: application/hal+json"</mgmt-ip>
```

The response:

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

```
### Deleting a volume efficiency policy
The DELETE operation is used to delete a volume efficiency policy.
```

The API:

```
/api/storage/volume-efficiency-policies/{uuid}
```

The call:

```
curl -X DELETE "https://<mgmt-ip>/api/storage/volume-efficiency-policies/ ae9e65c4-4506-11e9-aa44-005056bbc848" -H "accept: application/hal+json"</mgmt-ip>
```

The response:

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

```
[[IDe03bc3dd723c3634a7cc1a51069ecb2e]]
= Retrieve volume efficiency policies

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/volume-efficiency-policies`#

*Introduced In:* 9.8

Retrieves a collection of volume efficiency policies.

== Related ONTAP commands

* `volume efficiency policy show`

== Learn more

* xref:{relative_path}storage_volume-efficiency-
policies_endpoint_overview.html[DOC /storage/volume-efficiency-policies]

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
```



```
|Description

|enabled
|boolean
|query
|False
a|Filter by enabled

|qos_policy
|string
|query
|False
a|Filter by qos_policy

|uuid
|string
|query
|False
a|Filter by uuid

|comment
|string
|query
|False
a|Filter by comment

|name
|string
|query
|False
a|Filter by name

|type
|string
|query
|False
a|Filter by type

|duration
|integer
|query
```

```
|False
a|Filter by duration

|schedule.name
|string
|query
|False
a|Filter by schedule.name

|svm.uuid
|string
|query
|False
a|Filter by svm.uuid

|svm.name
|string
|query
|False
a|Filter by svm.name

|start_threshold_percent
|integer
|query
|False
a|Filter by start_threshold_percent

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
```

```
|boolean
|query
|False
a|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
a|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
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|num_records
|integer
a|Number of records

|records
|array[link:#volume_efficiency_policy[volume_efficiency_policy]]
a|

|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "comment": "string",
    "duration": 5,
    "enabled": 1,
    "name": "default",
    "qos_policy": "background",
    "schedule": {
      "name": "daily"
    },
    "start_threshold_percent": 30,
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  }
}
```

```

    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "threshold",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

```

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#schedule]
[.api-collapsible-fifth-title]
schedule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Schedule at which volume efficiency policies are captured on the SVM.
Some common schedules already defined in the system are hourly, daily,
weekly, at 5 minute intervals, and at 8 hour intervals. Volume efficiency
policies with custom schedules can be referenced.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#volume_efficiency_policy]
[.api-collapsible-fifth-title]
volume_efficiency_policy
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|comment
|string
a|A comment associated with the volume efficiency policy.
```

```
|duration
|integer
a|This field is used with the policy type "scheduled" to indicate the
allowed duration for a session, in hours. Possible value is a number
between 0 and 999 inclusive. Default is unlimited indicated by value 0.
```

```
|enabled
|boolean
a|Is the volume efficiency policy enabled?
```

```
|name
|string
```


a|Name of the volume efficiency policy.

|qos_policy

|string

a|QoS policy for the sis operation. Possible values are background and best_effort. In background, sis operation will run in background with minimal or no impact on data serving client operations. In best_effort, sis operations may have some impact on data serving client operations.

|schedule

|link:#schedule[schedule]

a|

|start_threshold_percent

|integer

a|This field is used with the policy type "threshold" to indicate the threshold percentage for triggering the volume efficiency policy. It is mutually exclusive of the schedule.

|svm

|link:#svm[svm]

a|

|type

|string

a|Type of volume efficiency policy.

|uuid

|string

a|Unique identifier of volume efficiency policy.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[ID434bd13b2b432f17d4094f973c505691]]
```

```
= Create a volume efficiency policy
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-  
block]#`/storage/volume-efficiency-policies`#
```

```
*Introduced In:* 9.8
```

```
Creates a volume efficiency policy.
```

```
== Required properties
```

```
* `svm.uuid` or `svm.name` - Existing SVM in which to create the volume  
efficiency policy.
```

```
* `name` - Name for the volume efficiency policy.
```

```
== Recommended optional properties
```

```
* `type` - Type of volume policy.
```

```
* `schedule` - Schedule the volume efficiency defined in minutes, hourly,  
daily and weekly.
```

```
* `duration` - Indicates the allowed duration for a session for policy  
type "scheduled".
```

```
* `start_threshold_percent` - Indicates the start threshold percentage for  
the policy type "threshold". It is mutually exclusive of the schedule.
```

```
* `qos_policy` - QoS policy for the sis operation.
```

```
* `comment` - A comment associated with the volume efficiency policy.
```

```
* `enabled` - Is the volume efficiency policy enabled?
```

```
== Default property values
```

```
If not specified in POST, the following default property values are  
assigned:
```

```
* `type` - scheduled
```

```
* `enabled` - true
```

```
* `qos_policy` - best_effort
```

```
== Related ONTAP commands
```

```
* `volume efficiency policy create`
```

== Learn more

* xref:{relative_path}storage_volume-efficiency-policies_endpoint_overview.html[DOC /storage/volume-efficiency-policies]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|comment

|string

a|A comment associated with the volume efficiency policy.

```
|duration
|integer
a|This field is used with the policy type "scheduled" to indicate the
allowed duration for a session, in hours. Possible value is a number
between 0 and 999 inclusive. Default is unlimited indicated by value 0.

|enabled
|boolean
a|Is the volume efficiency policy enabled?

|name
|string
a|Name of the volume efficiency policy.

|qos_policy
|string
a|QoS policy for the sis operation. Possible values are background and
best_effort. In background, sis operation will run in background with
minimal or no impact on data serving client operations. In best_effort,
sis operations may have some impact on data serving client operations.

|schedule
|link:#schedule[schedule]
a|

|start_threshold_percent
|integer
a|This field is used with the policy type "threshold" to indicate the
threshold percentage for triggering the volume efficiency policy. It is
mutually exclusive of the schedule.

|svm
|link:#svm[svm]
a|

|type
|string
a|Type of volume efficiency policy.

|uuid
|string
```

a|Unique identifier of volume efficiency policy.

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "duration": 5,
  "enabled": 1,
  "name": "default",
  "qos_policy": "background",
  "schedule": {
    "name": "daily"
  },
  "start_threshold_percent": 30,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "threshold",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

====

== Response

Status: 201, Created

```

=== Headers

[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 6881341
| Specified schedule not found.

| 6881344
| Failed to queue specified job.

| 6881345
| This operation is not permitted on a node SVM.

| 6881349
| Policy name is not valid.

| 6881362
| Threshold percentage cannot be less than 1 percent.

| 6881433
| For "{\{0\}}" type policy, attribute "{\{1\}}" is not supported.

```

```
| 6881435
| Only a policy of type "threshold" can set the "start-threshold-percent"
attribute.

| 6881436
| For a policy of type "scheduled", a valid "schedule" is a required
attribute.

| 6881454
| An efficiency policy of type "threshold" requires an effective cluster
version of ONTAP 8.3 or later.

| 6881474
| Duration cannot be null.

| 6881475
| Duration is not valid.

| 6881476
| Duration cannot be less than 1 hour.

| 6881477
| Duration cannot be more than 999 hours.
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```
.Example error
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{
```

```
  "error": {
```

```
    "arguments": {
```



```

    "code": "string",
    "message": "string"
  },
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}
}
====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]

```

```

a|

|===

[#schedule]
[.api-collapsible-fifth-title]
schedule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Schedule at which volume efficiency policies are captured on the SVM.
Some common schedules already defined in the system are hourly, daily,
weekly, at 5 minute intervals, and at 8 hour intervals. Volume efficiency
policies with custom schedules can be referenced.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string

```

```

a|The unique identifier of the SVM.

|===

[#volume_efficiency_policy]
[.api-collapsible-fifth-title]
volume_efficiency_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|comment
|string
a|A comment associated with the volume efficiency policy.

|duration
|integer
a|This field is used with the policy type "scheduled" to indicate the
allowed duration for a session, in hours. Possible value is a number
between 0 and 999 inclusive. Default is unlimited indicated by value 0.

|enabled
|boolean
a|Is the volume efficiency policy enabled?

|name
|string
a|Name of the volume efficiency policy.

|qos_policy
|string
a|QoS policy for the sis operation. Possible values are background and
best_effort. In background, sis operation will run in background with
minimal or no impact on data serving client operations. In best_effort,

```

sis operations may have some impact on data serving client operations.

|schedule
|link:#schedule[schedule]
a|

|start_threshold_percent
|integer
a|This field is used with the policy type "threshold" to indicate the threshold percentage for triggering the volume efficiency policy. It is mutually exclusive of the schedule.

|svm
|link:#svm[svm]
a|

|type
|string
a|Type of volume efficiency policy.

|uuid
|string
a|Unique identifier of volume efficiency policy.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message

```

|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID3fab233394d2d77d20d74740ecbdc647]]
= Delete a volume efficiency policy

```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/storage/volume-efficiency-policies/{uuid}`#
```

Introduced In: 9.8

Deletes a volume efficiency policy.

== Related ONTAP commands

* `volume efficiency policy modify`

== Learn more

* xref:{relative_path}storage_volume-efficiency-policies_endpoint_overview.html[DOC /storage/volume-efficiency-policies]

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Volume efficiency policy UUID
```

```
|===
```

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Code

|===

| Error Code | Description

| 6881346

| The policy was not deleted because the policy is in use by at least one volume.

| 6881347

| This operation cannot be performed because the specified policy is owned by the cluster admin.

| 6881431

| The specified policy is a predefined policy and cannot be deleted.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

====

[source,json,subs=+macros]

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

```

====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

```



```

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID1dc910362bbefba6f2c26fa2de2229da]]
= Retrieve volume efficiency policy details

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/volume-efficiency-policies/{uuid}`#

*Introduced In:* 9.8

Retrieves the details of the specified volume efficiency policy.

== Related ONTAP commands

* `volume efficiency policy show`

== Learn more

* xref:{relative_path}storage_volume-efficiency-
policies_endpoint_overview.html[DOC /storage/volume-efficiency-policies]

== Parameters

[cols=5*,options=header]
|===

```

```
|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Volume efficiency policy UUID

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|comment
|string
a|A comment associated with the volume efficiency policy.

|duration
|integer
a|This field is used with the policy type "scheduled" to indicate the
allowed duration for a session, in hours. Possible value is a number
between 0 and 999 inclusive. Default is unlimited indicated by value 0.
```

|enabled
|boolean
a|Is the volume efficiency policy enabled?

|name
|string
a|Name of the volume efficiency policy.

|qos_policy
|string
a|QoS policy for the sis operation. Possible values are background and best_effort. In background, sis operation will run in background with minimal or no impact on data serving client operations. In best_effort, sis operations may have some impact on data serving client operations.

|schedule
|link:#schedule[schedule]
a|

|start_threshold_percent
|integer
a|This field is used with the policy type "threshold" to indicate the threshold percentage for triggering the volume efficiency policy. It is mutually exclusive of the schedule.

|svm
|link:#svm[svm]
a|

|type
|string
a|Type of volume efficiency policy.

|uuid
|string
a|Unique identifier of volume efficiency policy.

|===

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "duration": 5,
  "enabled": 1,
  "name": "default",
  "qos_policy": "background",
  "schedule": {
    "name": "daily"
  },
  "start_threshold_percent": 30,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "threshold",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
```

```

|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#schedule]
[.api-collapsible-fifth-title]
schedule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Schedule at which volume efficiency policies are captured on the SVM.
Some common schedules already defined in the system are hourly, daily,
weekly, at 5 minute intervals, and at 8 hour intervals. Volume efficiency
policies with custom schedules can be referenced.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDd353360cb3867c4d07010d4530662b64]]
= Update a volume efficiency policy

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/volume-efficiency-policies/{uuid}`#

*Introduced In:* 9.8

Updates a volume efficiency policy.

== Related ONTAP commands

* `volume efficiency policy modify`

== Learn more

```



```
* xref:{relative_path}storage_volume-efficiency-
policies_endpoint_overview.html[DOC /storage/volume-efficiency-policies]
```

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Volume efficiency policy UUID
```

```
|===
```

== Request Body

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
 |_links
```

```
 |link:#_links[_links]
```

```
a|
```

```
 |comment
```

```
 |string
```

```
a|A comment associated with the volume efficiency policy.
```

```
 |duration
```

```
 |integer
```

```
a|This field is used with the policy type "scheduled" to indicate the
allowed duration for a session, in hours. Possible value is a number
```

between 0 and 999 inclusive. Default is unlimited indicated by value 0.

|enabled

|boolean

a|Is the volume efficiency policy enabled?

|name

|string

a|Name of the volume efficiency policy.

|qos_policy

|string

a|QoS policy for the sis operation. Possible values are background and best_effort. In background, sis operation will run in background with minimal or no impact on data serving client operations. In best_effort, sis operations may have some impact on data serving client operations.

|schedule

|link:#schedule[schedule]

a|

|start_threshold_percent

|integer

a|This field is used with the policy type "threshold" to indicate the threshold percentage for triggering the volume efficiency policy. It is mutually exclusive of the schedule.

|svm

|link:#svm[svm]

a|

|type

|string

a|Type of volume efficiency policy.

|uuid

|string

a|Unique identifier of volume efficiency policy.

|===

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "duration": 5,
  "enabled": 1,
  "name": "default",
  "qos_policy": "background",
  "schedule": {
    "name": "daily"
  },
  "start_threshold_percent": 30,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "threshold",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Code
```

```
|===  
| Error Code | Description  
  
| 6881341  
| Specified schedule not found.  
  
| 6881344  
| Failed to queue specified job.  
  
| 6881348  
| This operation cannot be performed because the specified policy is owned  
by the cluster admin.  
  
| 6881349  
| Policy name is not valid.  
  
| 6881362  
| Threshold percentage cannot be less than 1 percent.  
  
| 6881433  
| For "{\{0\}}" type policy, "{\{1\}}" duration is not supported.  
  
| 6881435  
| Only a policy of type "threshold" can set the "start-threshold-percent"  
attribute.  
  
| 6881436  
| For a policy of type "scheduled", a valid "schedule" is a required  
attribute.  
  
| 6881438  
| For "{\{0\}}" policy, modification of attributes is not allowed.  
  
| 6881474  
| Duration cannot be null.  
  
| 6881475  
| Duration is not valid.  
  
| 6881476  
| Duration cannot be less than 1 hour.  
  
| 6881477  
| Duration cannot be more than 999 hours.  
|===
```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#schedule]
[.api-collapsible-fifth-title]
schedule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Schedule at which volume efficiency policies are captured on the SVM.
Some common schedules already defined in the system are hourly, daily,
weekly, at 5 minute intervals, and at 8 hour intervals. Volume efficiency
policies with custom schedules can be referenced.

|===

```

```

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#volume_efficiency_policy]
[.api-collapsible-fifth-title]
volume_efficiency_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|comment
|string
a|A comment associated with the volume efficiency policy.

|duration

```

```
|integer
a|This field is used with the policy type "scheduled" to indicate the
allowed duration for a session, in hours. Possible value is a number
between 0 and 999 inclusive. Default is unlimited indicated by value 0.

|enabled
|boolean
a|Is the volume efficiency policy enabled?

|name
|string
a|Name of the volume efficiency policy.

|qos_policy
|string
a|QoS policy for the sis operation. Possible values are background and
best_effort. In background, sis operation will run in background with
minimal or no impact on data serving client operations. In best_effort,
sis operations may have some impact on data serving client operations.

|schedule
|link:#schedule[schedule]
a|

|start_threshold_percent
|integer
a|This field is used with the policy type "threshold" to indicate the
threshold percentage for triggering the volume efficiency policy. It is
mutually exclusive of the schedule.

|svm
|link:#svm[svm]
a|

|type
|string
a|Type of volume efficiency policy.

|uuid
|string
a|Unique identifier of volume efficiency policy.
```



```
|===
```

```
[#error_arguments]  
[.api-collapsible-fifth-title]  
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|code  
|string  
a|Argument code
```

```
|message  
|string  
a|Message argument
```

```
|===
```

```
[#error]  
[.api-collapsible-fifth-title]  
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|arguments  
|array[link:#error_arguments[error_arguments]]  
a|Message arguments
```

```
|code  
|string  
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
:leveloffset: -1
```

```
= Manage storage volumes
```

```
:leveloffset: +1
```

```
[[ID00595e7482bba485cb0813db64d19905]]
= 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:

```
&ndash; PATCH /api/storage/volumes/{uuid} -d '{"quota.enabled":"true"}
```

```
&ndash; PATCH /api/storage/volumes/{uuid} -d '{"quota.enabled":"false"}
```

```
&ndash; GET /api/storage/volumes/{uuid}/?fields=quota.state
```

== File System Analytics

File system analytics provide a quick method for obtaining information summarizing properties of all files within any directory tree of a volume. For more information on file system analytics, see [xref:{relative_path}storage_volumes_volume.uuid_files_path_endpoint_overview.html](#) [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:

```
&ndash; PATCH /api/storage/volumes/{uuid} -d '{"analytics.state":"on"}'
```

```
&ndash; PATCH /api/storage/volumes/{uuid} -d '{"analytics.state":"off"}'
```

```
&ndash; GET /api/storage/volumes/{uuid}/?fields=analytics
```

== QoS

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

== Performance monitoring

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

== 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 rebalancing 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:

```
&ndash; POST      /api/storage/volumes
&ndash; GET      /api/storage/volumes
&ndash; GET      /api/storage/volumes/{uuid}
&ndash; PATCH    /api/storage/volumes/{uuid}
&ndash; DELETE   /api/storage/volumes/{uuid}
```

== Examples

=== Creating a volume

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

The API:

```
/api/storage/volumes
```

The call:

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

The response:

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

=== Creating a 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": "voll", "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": "voll_clone", "clone": {"parent_volume":
{"name": "voll"}, "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": "voll",
  "uuid": "a2564f80-25fb-41e8-9b49-44de2600991f",
  "_links": {
    "self": {
```

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

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

```

    "status": "ok",
    "cache_miss_blocks": 0,
    "client_requested_blocks": 0
  }
},
"qos": {
  "policy": {
    "min_throughput_iops": 0,
    "min_throughput_mbps": 0,
    "max_throughput_iops": 1000,
    "max_throughput_mbps": 0,
    "uuid": "228454af-5a8b-11e9-bd5b-005056ac6f1f",
    "name": "pg1"
  }
},
"snaplock": {
  "append_mode_enabled": false,
  "autocommit_period": "none",
  "compliance_clock_time": "2019-05-24T10:59:00+05:30",
  "expiry_time": "2038-01-19T08:44:28+05:30",
  "is_audit_log": false,
  "litigation_count": 0,
  "privileged_delete": "disabled",
  "type": "enterprise",
  "retention": {
    "default": "P0Y",
    "minimum": "P0Y",
    "maximum": "P30Y"
  }
},
"snapshot_policy": {
  "name": "default"
},
"svm": {
  "name": "vsdata",
  "uuid": "d61b69f5-7458-11e8-ad3f-0050569503ac"
},
"anti_ransomware_state": "disabled",
"_links": {
  "self": {
    "href": "/api/storage/volumes/d0c3359c-5448-4a9b-a077-e3295a7e9057"
  }
}
}
}
-----

```

```
=== Retrieving the quota state of a FlexVol or a FlexGroup volume
```

```
----
```

```
# The API:
```

```
/api/storage/volumes/{uuid}
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717/?fields=quota.state" -H "accept: application/hal+json"
```

```
# The response:
```

```
{  
  "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",  
  "name": "fv",  
  "quota": {  
    "state": "on"  
  },  
  "_links": {  
    "self": {  
      "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717/"  
    }  
  }  
}
```

```
----
```

```
=== Retrieving the constituents of a FlexGroup volume
```

```
----
```

```
# The API:
```

```
/api/storage/volumes/{uuid}
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes?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_contituent=true&flexg
roup.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_contituent=true&flexgro
up.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"
  }
}
}
}
-----
```

```
[[IDb6c8e7f3d6d6d9355f223bc14621da3a]]
```

```
= Retrieve volumes
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/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 [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) 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.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`
```

```
* `tiering.storage_class`
```

```
== 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
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|is_constituent
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|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
```

```
|flexgroup.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by flexgroup.uuid
```

```
* Introduced in: 9.10
```

```
|flexgroup.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by flexgroup.name
```

```
* Introduced in: 9.10
```

```
* maxLength: 203
```

```
* minLength: 1
```

```
|svm.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by svm.uuid
```

```
|svm.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by svm.name
```

```
|nas.gid
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by nas.gid
```

```
|nas.unix_permissions
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by nas.unix_permissions
```

```
|nas.uid
```

```
|integer
```

```
|query
```

```
|False  
a|Filter by nas.uid
```

```
|nas.path  
|string  
|query  
|False  
a|Filter by nas.path
```

```
|nas.security_style  
|string  
|query  
|False  
a|Filter by nas.security_style
```

```
|nas.junction_parent.uuid  
|string  
|query  
|False  
a|Filter by nas.junction_parent.uuid
```

* Introduced in: 9.9

```
|nas.junction_parent.name  
|string  
|query  
|False  
a|Filter by nas.junction_parent.name
```

* Introduced in: 9.9

```
|nas.export_policy.name  
|string  
|query  
|False  
a|Filter by nas.export_policy.name
```

```
|nas.export_policy.id  
|integer  
|query  
|False
```

```
a|Filter by nas.export_policy.id
```

```
|space.cross_volume_dedupe_metafiles_temporary_footprint  
|integer  
|query  
|False
```

```
a|Filter by space.cross_volume_dedupe_metafiles_temporary_footprint
```

```
* Introduced in: 9.10
```

```
|space.user_data  
|integer  
|query  
|False
```

```
a|Filter by space.user_data
```

```
* Introduced in: 9.10
```

```
|space.snapshot.autodelete.commitment  
|string  
|query  
|False
```

```
a|Filter by space.snapshot.autodelete.commitment
```

```
* Introduced in: 9.13
```

```
|space.snapshot.autodelete.prefix  
|string  
|query  
|False
```

```
a|Filter by space.snapshot.autodelete.prefix
```

```
* Introduced in: 9.13
```

```
|space.snapshot.autodelete.target_free_space  
|integer  
|query  
|False
```

```
a|Filter by space.snapshot.autodelete.target_free_space
```

```
* Introduced in: 9.13
```

```
|space.snapshot.autodelete.defer_delete
|string
|query
|False
a|Filter by space.snapshot.autodelete.defer_delete
```

* Introduced in: 9.13

```
|space.snapshot.autodelete.delete_order
|string
|query
|False
a|Filter by space.snapshot.autodelete.delete_order
```

* Introduced in: 9.13

```
|space.snapshot.autodelete.trigger
|string
|query
|False
a|Filter by space.snapshot.autodelete.trigger
```

* Introduced in: 9.13

```
|space.snapshot.autodelete_trigger
|string
|query
|False
a|Filter by space.snapshot.autodelete_trigger
```

* Introduced in: 9.10

```
|space.snapshot.space_used_percent
|integer
|query
|False
a|Filter by space.snapshot.space_used_percent
```

* Introduced in: 9.9

```
|space.snapshot.reserve_size
```

```
|integer
|query
|False
a|Filter by space.snapshot.reserve_size

* Introduced in: 9.9

|space.snapshot.reserve_percent
|integer
|query
|False
a|Filter by space.snapshot.reserve_percent

|space.snapshot.used
|integer
|query
|False
a|Filter by space.snapshot.used

|space.snapshot.reserve_available
|integer
|query
|False
a|Filter by space.snapshot.reserve_available

* Introduced in: 9.10

|space.percent_used
|integer
|query
|False
a|Filter by space.percent_used

* Introduced in: 9.9

|space.overwrite_reserve_used
|integer
|query
|False
a|Filter by space.overwrite_reserve_used

* Introduced in: 9.9
```

```
|space.used
|integer
|query
|False
a|Filter by space.used
```

```
|space.is_used_stale
|boolean
|query
|False
a|Filter by space.is_used_stale
```

* Introduced in: 9.12

```
|space.large_size_enabled
|boolean
|query
|False
a|Filter by space.large_size_enabled
```

* Introduced in: 9.12

```
|space.dedupe_metafiles_temporary_footprint
|integer
|query
|False
a|Filter by space.dedupe_metafiles_temporary_footprint
```

* Introduced in: 9.10

```
|space.block_storage_inactive_user_data_percent
|integer
|query
|False
a|Filter by space.block_storage_inactive_user_data_percent
```

* Introduced in: 9.9

```
|space.used_by_afs
|integer
```



```
|query
|False
a|Filter by space.used_by_afs
```

* Introduced in: 9.9

```
|space.local_tier_footprint
|integer
|query
|False
a|Filter by space.local_tier_footprint
```

* Introduced in: 9.8

```
|space.expected_available
|integer
|query
|False
a|Filter by space.expected_available
```

* Introduced in: 9.10

```
|space.overwrite_reserve
|integer
|query
|False
a|Filter by space.overwrite_reserve
```

* Introduced in: 9.9

```
|space.physical_used
|integer
|query
|False
a|Filter by space.physical_used
```

* Introduced in: 9.10

```
|space.size_available_for_snapshots
|integer
|query
|False
```

a|Filter by space.size_available_for_snapshots

* Introduced in: 9.9

|space.available_percent

|integer

|query

|False

a|Filter by space.available_percent

* Introduced in: 9.9

|space.cross_volume_dedupe_metafiles_footprint

|integer

|query

|False

a|Filter by space.cross_volume_dedupe_metafiles_footprint

* Introduced in: 9.10

|space.size

|integer

|query

|False

a|Filter by space.size

|space.available

|integer

|query

|False

a|Filter by space.available

|space.total_footprint

|integer

|query

|False

a|Filter by space.total_footprint

* Introduced in: 9.8

|space.block_storage_inactive_user_data

```
|integer
|query
|False
a|Filter by space.block_storage_inactive_user_data
```

```
|space.capacity_tier_footprint_data_reduction
|integer
|query
|False
a|Filter by space.capacity_tier_footprint_data_reduction
```

* Introduced in: 9.13

```
|space.file_operation_metadata
|integer
|query
|False
a|Filter by space.file_operation_metadata
```

* Introduced in: 9.10

```
|space.over_provisioned
|integer
|query
|False
a|Filter by space.over_provisioned
```

```
|space.footprint
|integer
|query
|False
a|Filter by space.footprint
```

```
|space.snapshot_spill
|integer
|query
|False
a|Filter by space.snapshot_spill
```

* Introduced in: 9.10

```
|space.physical_used_percent
|integer
|query
|False
a|Filter by space.physical_used_percent
```

* Introduced in: 9.10

```
|space.dedupe_metafiles_footprint
|integer
|query
|False
a|Filter by space.dedupe_metafiles_footprint
```

* Introduced in: 9.10

```
|space.nearly_full_threshold_percent
|integer
|query
|False
a|Filter by space.nearly_full_threshold_percent
```

* Introduced in: 9.9

```
|space.filesystem_size
|integer
|query
|False
a|Filter by space.filesystem_size
```

* Introduced in: 9.10

```
|space.snapshot_reserve_unusable
|integer
|query
|False
a|Filter by space.snapshot_reserve_unusable
```

* Introduced in: 9.10

```
|space.afs_total
|integer
```

```
|query
|False
a|Filter by space.afs_total
```

* Introduced in: 9.9

```
|space.delayed_free_footprint
|integer
|query
|False
a|Filter by space.delayed_free_footprint
```

* Introduced in: 9.10

```
|space.auto_adaptive_compression_footprint_data_reduction
|integer
|query
|False
a|Filter by space.auto_adaptive_compression_footprint_data_reduction
```

* Introduced in: 9.11

```
|space.fractional_reserve
|integer
|query
|False
a|Filter by space.fractional_reserve
```

* Introduced in: 9.9

```
|space.volume_guarantee_footprint
|integer
|query
|False
a|Filter by space.volume_guarantee_footprint
```

* Introduced in: 9.10

```
|space.full_threshold_percent
|integer
|query
|False
```

a|Filter by space.full_threshold_percent

* Introduced in: 9.9

|space.filesystem_size_fixed

|boolean

|query

|False

a|Filter by space.filesystem_size_fixed

* Introduced in: 9.10

|space.metadata

|integer

|query

|False

a|Filter by space.metadata

|space.snapmirror_destination_footprint

|integer

|query

|False

a|Filter by space.snapmirror_destination_footprint

* Introduced in: 9.10

|space.logical_space.used_percent

|integer

|query

|False

a|Filter by space.logical_space.used_percent

* Introduced in: 9.9

|space.logical_space.reporting

|boolean

|query

|False

a|Filter by space.logical_space.reporting

|space.logical_space.used

```
|integer
|query
|False
a|Filter by space.logical_space.used

* Introduced in: 9.9

|space.logical_space.available
|integer
|query
|False
a|Filter by space.logical_space.available

|space.logical_space.enforcement
|boolean
|query
|False
a|Filter by space.logical_space.enforcement

|space.logical_space.used_by_snapshots
|integer
|query
|False
a|Filter by space.logical_space.used_by_snapshots

* Introduced in: 9.10

|space.logical_space.used_by_afs
|integer
|query
|False
a|Filter by space.logical_space.used_by_afs

|space.effective_total_footprint
|integer
|query
|False
a|Filter by space.effective_total_footprint

* Introduced in: 9.11
```

```
|space.capacity_tier_footprint
|integer
|query
|False
a|Filter by space.capacity_tier_footprint

|space.performance_tier_footprint
|integer
|query
|False
a|Filter by space.performance_tier_footprint

* Introduced in: 9.8

|snaplock.append_mode_enabled
|boolean
|query
|False
a|Filter by snaplock.append_mode_enabled

|snaplock.retention.minimum
|string
|query
|False
a|Filter by snaplock.retention.minimum

|snaplock.retention.default
|string
|query
|False
a|Filter by snaplock.retention.default

|snaplock.retention.maximum
|string
|query
|False
a|Filter by snaplock.retention.maximum

|snaplock.compliance_clock_time
|string
|query
```



```
|False
a|Filter by snaplock.compliance_clock_time

|snaplock.expiry_time
|string
|query
|False
a|Filter by snaplock.expiry_time

|snaplock.litigation_count
|integer
|query
|False
a|Filter by snaplock.litigation_count

|snaplock.is_audit_log
|boolean
|query
|False
a|Filter by snaplock.is_audit_log

|snaplock.privileged_delete
|string
|query
|False
a|Filter by snaplock.privileged_delete

|snaplock.type
|string
|query
|False
a|Filter by snaplock.type

|snaplock.unspecified_retention_file_count
|integer
|query
|False
a|Filter by snaplock.unspecified_retention_file_count

* Introduced in: 9.8
```

```
|snaplock.autocommit_period
|string
|query
|False
a|Filter by snaplock.autocommit_period
```

```
|efficiency.policy.name
|string
|query
|False
a|Filter by efficiency.policy.name
```

* Introduced in: 9.7

```
|efficiency.type
|string
|query
|False
a|Filter by efficiency.type
```

* Introduced in: 9.9

```
|efficiency.scanner.compression
|boolean
|query
|False
a|Filter by efficiency.scanner.compression
```

* Introduced in: 9.11

```
|efficiency.scanner.dedupe
|boolean
|query
|False
a|Filter by efficiency.scanner.dedupe
```

* Introduced in: 9.11

```
|efficiency.scanner.state
|string
|query
```

```
|False
a|Filter by efficiency.scanner.state

* Introduced in: 9.11

|efficiency.scanner.scan_old_data
|boolean
|query
|False
a|Filter by efficiency.scanner.scan_old_data

* Introduced in: 9.11

|efficiency.schedule
|string
|query
|False
a|Filter by efficiency.schedule

* Introduced in: 9.8

|efficiency.last_op_begin
|string
|query
|False
a|Filter by efficiency.last_op_begin

* Introduced in: 9.9

|efficiency.storage_efficiency_mode
|string
|query
|False
a|Filter by efficiency.storage_efficiency_mode

* Introduced in: 9.10

|efficiency.compression
|string
|query
|False
a|Filter by efficiency.compression
```

```
|efficiency.progress
|string
|query
|False
a|Filter by efficiency.progress
```

* Introduced in: 9.9

```
|efficiency.has_savings
|boolean
|query
|False
a|Filter by efficiency.has_savings
```

* Introduced in: 9.11

```
|efficiency.volume_path
|string
|query
|False
a|Filter by efficiency.volume_path
```

* Introduced in: 9.13

```
|efficiency.application_io_size
|string
|query
|False
a|Filter by efficiency.application_io_size
```

* Introduced in: 9.8

```
|efficiency.auto_state
|string
|query
|False
a|Filter by efficiency.auto_state
```

* Introduced in: 9.12

```
|efficiency.cross_volume_dedupe
|string
|query
|False
a|Filter by efficiency.cross_volume_dedupe
```

```
|efficiency.state
|string
|query
|False
a|Filter by efficiency.state
```

* Introduced in: 9.9

```
|efficiency.compaction
|string
|query
|False
a|Filter by efficiency.compaction
```

```
|efficiency.op_state
|string
|query
|False
a|Filter by efficiency.op_state
```

* Introduced in: 9.9

```
|efficiency.last_op_size
|integer
|query
|False
a|Filter by efficiency.last_op_size
```

* Introduced in: 9.9

```
|efficiency.dedupe
|string
|query
|False
a|Filter by efficiency.dedupe
```

```
|efficiency.space_savings.compression
|integer
|query
|False
a|Filter by efficiency.space_savings.compression
```

* Introduced in: 9.11

```
|efficiency.space_savings.total
|integer
|query
|False
a|Filter by efficiency.space_savings.total
```

* Introduced in: 9.11

```
|efficiency.space_savings.dedupe_percent
|integer
|query
|False
a|Filter by efficiency.space_savings.dedupe_percent
```

* Introduced in: 9.11

```
|efficiency.space_savings.total_percent
|integer
|query
|False
a|Filter by efficiency.space_savings.total_percent
```

* Introduced in: 9.11

```
|efficiency.space_savings.dedupe
|integer
|query
|False
a|Filter by efficiency.space_savings.dedupe
```

* Introduced in: 9.11

```
|efficiency.space_savings.compression_percent
```

```
|integer
|query
|False
a|Filter by efficiency.space_savings.compression_percent
```

* Introduced in: 9.11

```
|efficiency.space_savings.dedupe_sharing
|integer
|query
|False
a|Filter by efficiency.space_savings.dedupe_sharing
```

* Introduced in: 9.11

```
|efficiency.logging_enabled
|boolean
|query
|False
a|Filter by efficiency.logging_enabled
```

* Introduced in: 9.11

```
|efficiency.last_op_state
|string
|query
|False
a|Filter by efficiency.last_op_state
```

* Introduced in: 9.9

```
|efficiency.compression_type
|string
|query
|False
a|Filter by efficiency.compression_type
```

* Introduced in: 9.11

```
|efficiency.last_op_end
|string
|query
```

```
|False
a|Filter by efficiency.last_op_end

* Introduced in: 9.9

|efficiency.idcs_scanner.status
|string
|query
|False
a|Filter by efficiency.idcs_scanner.status

* Introduced in: 9.13

|efficiency.idcs_scanner.enabled
|boolean
|query
|False
a|Filter by efficiency.idcs_scanner.enabled

* Introduced in: 9.13

|efficiency.idcs_scanner.threshold_inactive_time
|string
|query
|False
a|Filter by efficiency.idcs_scanner.threshold_inactive_time

* Introduced in: 9.13

|efficiency.idcs_scanner.mode
|string
|query
|False
a|Filter by efficiency.idcs_scanner.mode

* Introduced in: 9.13

|efficiency.idcs_scanner.operation_state
|string
|query
|False
a|Filter by efficiency.idcs_scanner.operation_state
```


* Introduced in: 9.13

|efficiency.last_op_err
|string
|query
|False
a|Filter by efficiency.last_op_err

* Introduced in: 9.9

|granular_data
|boolean
|query
|False
a|Filter by granular_data

* Introduced in: 9.11

|is_svm_root
|boolean
|query
|False
a|Filter by is_svm_root

* Introduced in: 9.7

|autosize.grow_threshold
|integer
|query
|False
a|Filter by autosize.grow_threshold

|autosize.minimum
|integer
|query
|False
a|Filter by autosize.minimum

|autosize.maximum
|integer

```
|query
|False
a|Filter by autosize.maximum
```

```
|autosize.mode
|string
|query
|False
a|Filter by autosize.mode
```

```
|autosize.shrink_threshold
|integer
|query
|False
a|Filter by autosize.shrink_threshold
```

```
|aggregates.name
|string
|query
|False
a|Filter by aggregates.name
```

```
|aggregates.uuid
|string
|query
|False
a|Filter by aggregates.uuid
```

```
|cloud_retrieval_policy
|string
|query
|False
a|Filter by cloud_retrieval_policy
```

* Introduced in: 9.8

```
|consistency_group.uuid
|string
|query
|False
a|Filter by consistency_group.uuid
```

* Introduced in: 9.10

|consistency_group.name
|string
|query
|False
a|Filter by consistency_group.name

* Introduced in: 9.7

|snapshot_policy.uuid
|string
|query
|False
a|Filter by snapshot_policy.uuid

|snapshot_policy.name
|string
|query
|False
a|Filter by snapshot_policy.name

|comment
|string
|query
|False
a|Filter by comment

* maxLength: 1023

* minLength: 0

|language
|string
|query
|False
a|Filter by language

|snapshot_directory_access_enabled
|boolean
|query

```
|False  
a|Filter by snapshot_directory_access_enabled
```

```
* Introduced in: 9.13
```

```
|clone.split_initiated  
|boolean  
|query  
|False  
a|Filter by clone.split_initiated
```

```
|clone.parent_svm.uuid  
|string  
|query  
|False  
a|Filter by clone.parent_svm.uuid
```

```
|clone.parent_svm.name  
|string  
|query  
|False  
a|Filter by clone.parent_svm.name
```

```
|clone.inherited_physical_used  
|integer  
|query  
|False  
a|Filter by clone.inherited_physical_used
```

```
* Introduced in: 9.12
```

```
|clone.split_estimate  
|integer  
|query  
|False  
a|Filter by clone.split_estimate
```

```
|clone.split_complete_percent  
|integer  
|query  
|False
```

```
a|Filter by clone.split_complete_percent
```

```
|clone.is_flexclone
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|Filter by clone.is_flexclone
```

```
|clone.parent_snapshot.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by clone.parent_snapshot.name
```

```
|clone.parent_snapshot.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by clone.parent_snapshot.uuid
```

```
|clone.parent_volume.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by clone.parent_volume.uuid
```

```
|clone.parent_volume.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by clone.parent_volume.name
```

```
|clone.inherited_savings
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by clone.inherited_savings
```

```
* Introduced in: 9.12
```

```
|type
|string
|query
|False
a|Filter by type
```

```
|anti_ransomware.space.used
|integer
|query
|False
a|Filter by anti_ransomware.space.used
```

* Introduced in: 9.10

```
|anti_ransomware.space.used_by_logs
|integer
|query
|False
a|Filter by anti_ransomware.space.used_by_logs
```

* Introduced in: 9.10

```
|anti_ransomware.space.used_by_snapshots
|integer
|query
|False
a|Filter by anti_ransomware.space.used_by_snapshots
```

* Introduced in: 9.10

```
|anti_ransomware.space.snapshot_count
|integer
|query
|False
a|Filter by anti_ransomware.space.snapshot_count
```

* Introduced in: 9.10

```
|anti_ransomware.state
|string
|query
|False
```

a|Filter by anti_ransomware.state

* Introduced in: 9.10

|anti_ransomware.attack_probability

|string

|query

|False

a|Filter by anti_ransomware.attack_probability

* Introduced in: 9.10

|anti_ransomware.attack_reports.time

|string

|query

|False

a|Filter by anti_ransomware.attack_reports.time

* Introduced in: 9.10

|anti_ransomware.surge_as_normal

|boolean

|query

|False

a|Filter by anti_ransomware.surge_as_normal

* Introduced in: 9.11

|anti_ransomware.dry_run_start_time

|string

|query

|False

a|Filter by anti_ransomware.dry_run_start_time

* Introduced in: 9.10

|anti_ransomware.suspect_files.format

|string

|query

|False

a|Filter by anti_ransomware.suspect_files.format

* Introduced in: 9.10

|anti_ransomware.suspect_files.count
|integer
|query
|False
a|Filter by anti_ransomware.suspect_files.count

* Introduced in: 9.10

|anti_ransomware.suspect_files.entropy
|string
|query
|False
a|Filter by anti_ransomware.suspect_files.entropy

* Introduced in: 9.11

|quota.state
|string
|query
|False
a|Filter by quota.state

|queue_for_encryption
|boolean
|query
|False
a|Filter by queue_for_encryption

* Introduced in: 9.8

|asynchronous_directory_delete.enabled
|boolean
|query
|False
a|Filter by asynchronous_directory_delete.enabled

* Introduced in: 9.11

|asynchronous_directory_delete.trash_bin


```
|string
|query
|False
a|Filter by asynchronous_directory_delete.trash_bin
```

* Introduced in: 9.11

```
|activity_tracking.unsupported_reason.message
|string
|query
|False
a|Filter by activity_tracking.unsupported_reason.message
```

* Introduced in: 9.10

```
|activity_tracking.unsupported_reason.code
|string
|query
|False
a|Filter by activity_tracking.unsupported_reason.code
```

* Introduced in: 9.10

```
|activity_tracking.supported
|boolean
|query
|False
a|Filter by activity_tracking.supported
```

* Introduced in: 9.10

```
|activity_tracking.state
|string
|query
|False
a|Filter by activity_tracking.state
```

* Introduced in: 9.10

```
|flash_pool.cache_eligibility
|string
|query
```

```
|False
a|Filter by flash_pool.cache_eligibility

* Introduced in: 9.10

|flash_pool.caching_policy
|string
|query
|False
a|Filter by flash_pool.caching_policy

* Introduced in: 9.10

|flash_pool.cache_retention_priority
|string
|query
|False
a|Filter by flash_pool.cache_retention_priority

* Introduced in: 9.10

|aggressive_readahead_mode
|string
|query
|False
a|Filter by aggressive_readahead_mode

* Introduced in: 9.13

|access_time_enabled
|boolean
|query
|False
a|Filter by access_time_enabled

* Introduced in: 9.8

|error_state.is_inconsistent
|boolean
|query
|False
a|Filter by error_state.is_inconsistent
```

```
|error_state.has_bad_blocks
|boolean
|query
|False
a|Filter by error_state.has_bad_blocks
```

```
|scheduled_snapshot_naming_scheme
|string
|query
|False
a|Filter by scheduled_snapshot_naming_scheme
```

* Introduced in: 9.10

```
|guarantee.honored
|boolean
|query
|False
a|Filter by guarantee.honored
```

```
|guarantee.type
|string
|query
|False
a|Filter by guarantee.type
```

```
|application.uuid
|string
|query
|False
a|Filter by application.uuid
```

```
|application.name
|string
|query
|False
a|Filter by application.name
```

```
|create_time
```

```
|string  
|query  
|False  
a|Filter by create_time
```

```
|metric.latency.read  
|integer  
|query  
|False  
a|Filter by metric.latency.read
```

```
|metric.latency.other  
|integer  
|query  
|False  
a|Filter by metric.latency.other
```

```
|metric.latency.write  
|integer  
|query  
|False  
a|Filter by metric.latency.write
```

```
|metric.latency.total  
|integer  
|query  
|False  
a|Filter by metric.latency.total
```

```
|metric.throughput.read  
|integer  
|query  
|False  
a|Filter by metric.throughput.read
```

```
|metric.throughput.other  
|integer  
|query  
|False  
a|Filter by metric.throughput.other
```

```
|metric.throughput.write
|integer
|query
|False
a|Filter by metric.throughput.write
```

```
|metric.throughput.total
|integer
|query
|False
a|Filter by metric.throughput.total
```

```
|metric.duration
|string
|query
|False
a|Filter by metric.duration
```

```
|metric.flexcache.status
|string
|query
|False
a|Filter by metric.flexcache.status
```

* Introduced in: 9.8

```
|metric.flexcache.timestamp
|string
|query
|False
a|Filter by metric.flexcache.timestamp
```

* Introduced in: 9.8

```
|metric.flexcache.bandwidth_savings
|integer
|query
|False
a|Filter by metric.flexcache.bandwidth_savings
```

* Introduced in: 9.9

```
|metric.flexcache.duration  
|string  
|query  
|False  
a|Filter by metric.flexcache.duration
```

* Introduced in: 9.8

```
|metric.flexcache.cache_miss_percent  
|integer  
|query  
|False  
a|Filter by metric.flexcache.cache_miss_percent
```

* Introduced in: 9.8

```
|metric.iops.read  
|integer  
|query  
|False  
a|Filter by metric.iops.read
```

```
|metric.iops.other  
|integer  
|query  
|False  
a|Filter by metric.iops.other
```

```
|metric.iops.write  
|integer  
|query  
|False  
a|Filter by metric.iops.write
```

```
|metric.iops.total  
|integer  
|query  
|False  
a|Filter by metric.iops.total
```

```
|metric.cloud.iops.read  
|integer  
|query  
|False  
a|Filter by metric.cloud.iops.read
```

* Introduced in: 9.7

```
|metric.cloud.iops.other  
|integer  
|query  
|False  
a|Filter by metric.cloud.iops.other
```

* Introduced in: 9.7

```
|metric.cloud.iops.write  
|integer  
|query  
|False  
a|Filter by metric.cloud.iops.write
```

* Introduced in: 9.7

```
|metric.cloud.iops.total  
|integer  
|query  
|False  
a|Filter by metric.cloud.iops.total
```

* Introduced in: 9.7

```
|metric.cloud.latency.read  
|integer  
|query  
|False  
a|Filter by metric.cloud.latency.read
```

* Introduced in: 9.7

```
|metric.cloud.latency.other
```

```
|integer
|query
|False
a|Filter by metric.cloud.latency.other
```

* Introduced in: 9.7

```
|metric.cloud.latency.write
|integer
|query
|False
a|Filter by metric.cloud.latency.write
```

* Introduced in: 9.7

```
|metric.cloud.latency.total
|integer
|query
|False
a|Filter by metric.cloud.latency.total
```

* Introduced in: 9.7

```
|metric.cloud.status
|string
|query
|False
a|Filter by metric.cloud.status
```

* Introduced in: 9.7

```
|metric.cloud.timestamp
|string
|query
|False
a|Filter by metric.cloud.timestamp
```

* Introduced in: 9.7

```
|metric.cloud.duration
|string
|query
```



```
|False  
a|Filter by metric.cloud.duration
```

* Introduced in: 9.7

```
|metric.timestamp  
|string  
|query  
|False  
a|Filter by metric.timestamp
```

```
|metric.status  
|string  
|query  
|False  
a|Filter by metric.status
```

```
|encryption.rekey  
|boolean  
|query  
|False  
a|Filter by encryption.rekey
```

```
|encryption.key_create_time  
|string  
|query  
|False  
a|Filter by encryption.key_create_time
```

* Introduced in: 9.11

```
|encryption.type  
|string  
|query  
|False  
a|Filter by encryption.type
```

```
|encryption.state  
|string  
|query  
|False
```

a|Filter by encryption.state

|encryption.key_id

|string

|query

|False

a|Filter by encryption.key_id

|encryption.status.code

|string

|query

|False

a|Filter by encryption.status.code

|encryption.status.message

|string

|query

|False

a|Filter by encryption.status.message

|encryption.enabled

|boolean

|query

|False

a|Filter by encryption.enabled

|_tags

|string

|query

|False

a|Filter by _tags

* Introduced in: 9.13

|size

|integer

|query

|False

a|Filter by size

```
|statistics.cifs_ops_raw.link.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.link.total_time
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.link.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.link.count
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.unlink.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.unlink.total_time
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.unlink.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.unlink.count
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.lookup.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.lookup.total_time
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.lookup.count
|integer
```

```
|query
|False
a|Filter by statistics.cifs_ops_raw.lookup.count

* Introduced in: 9.11

|statistics.cifs_ops_raw.readlink.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.readlink.total_time

* Introduced in: 9.11

|statistics.cifs_ops_raw.readlink.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.readlink.count

* Introduced in: 9.11

|statistics.cifs_ops_raw.open.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.open.total_time

* Introduced in: 9.11

|statistics.cifs_ops_raw.open.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.open.count

* Introduced in: 9.11

|statistics.cifs_ops_raw.watch.total_time
|integer
|query
|False
```

a|Filter by statistics.cifs_ops_raw.watch.total_time

* Introduced in: 9.11

|statistics.cifs_ops_raw.watch.count

|integer

|query

|False

a|Filter by statistics.cifs_ops_raw.watch.count

* Introduced in: 9.11

|statistics.cifs_ops_raw.readdir.total_time

|integer

|query

|False

a|Filter by statistics.cifs_ops_raw.readdir.total_time

* Introduced in: 9.11

|statistics.cifs_ops_raw.readdir.count

|integer

|query

|False

a|Filter by statistics.cifs_ops_raw.readdir.count

* Introduced in: 9.11

|statistics.cifs_ops_raw.rename.total_time

|integer

|query

|False

a|Filter by statistics.cifs_ops_raw.rename.total_time

* Introduced in: 9.11

|statistics.cifs_ops_raw.rename.count

|integer

|query

|False

a|Filter by statistics.cifs_ops_raw.rename.count

* Introduced in: 9.11

|statistics.cifs_ops_raw.create.symlink.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.create.symlink.total_time

* Introduced in: 9.11

|statistics.cifs_ops_raw.create.symlink.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.create.symlink.count

* Introduced in: 9.11

|statistics.cifs_ops_raw.create.dir.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.create.dir.total_time

* Introduced in: 9.11

|statistics.cifs_ops_raw.create.dir.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.create.dir.count

* Introduced in: 9.11

|statistics.cifs_ops_raw.create.file.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.create.file.total_time

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.create.file.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.create.file.count
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.create.other.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.create.other.total_time
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.create.other.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.create.other.count
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.audit.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.audit.total_time
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.audit.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.audit.count
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.access.total_time
```

```
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.access.total_time
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.access.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.access.count
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.write.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.write.count
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.write.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.write.total_time
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.write.volume_protocol_latency_histogram_labels
|string
|query
|False
a|Filter by
statistics.cifs_ops_raw.write.volume_protocol_latency_histogram_labels
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.write.volume_protocol_size_histogram_labels
|string
```



```
|query
|False
a|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_counts
|integer
|query
|False
a|Filter by
statistics.cifs_ops_raw.write.volume_protocol_latency_histogram_counts

* Introduced in: 9.11

|statistics.cifs_ops_raw.write.volume_protocol_size_histogram_counts
|integer
|query
|False
a|Filter by
statistics.cifs_ops_raw.write.volume_protocol_size_histogram_counts

* Introduced in: 9.11

|statistics.cifs_ops_raw.getattr.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.getattr.total_time

* Introduced in: 9.11

|statistics.cifs_ops_raw.getattr.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.getattr.count

* Introduced in: 9.11

|statistics.cifs_ops_raw.read.count
```

```
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.read.count
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.read.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.read.total_time
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.read.volume_protocol_latency_histogram_labels
|string
|query
|False
a|Filter by
statistics.cifs_ops_raw.read.volume_protocol_latency_histogram_labels
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.read.volume_protocol_size_histogram_labels
|string
|query
|False
a|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_counts
|integer
|query
|False
a|Filter by
statistics.cifs_ops_raw.read.volume_protocol_latency_histogram_counts
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.read.volume_protocol_size_histogram_counts
|integer
|query
|False
a|Filter by
statistics.cifs_ops_raw.read.volume_protocol_size_histogram_counts
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.setattr.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.setattr.total_time
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.setattr.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.setattr.count
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.lock.total_time
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.lock.total_time
```

* Introduced in: 9.11

```
|statistics.cifs_ops_raw.lock.count
|integer
|query
|False
a|Filter by statistics.cifs_ops_raw.lock.count
```

* Introduced in: 9.11

```
|statistics.timestamp
```

```
|string
|query
|False
a|Filter by statistics.timestamp
```

```
|statistics.status
|string
|query
|False
a|Filter by statistics.status
```

```
|statistics.cloud.status
|string
|query
|False
a|Filter by statistics.cloud.status
```

* Introduced in: 9.7

```
|statistics.cloud.timestamp
|string
|query
|False
a|Filter by statistics.cloud.timestamp
```

* Introduced in: 9.7

```
|statistics.cloud.iops_raw.read
|integer
|query
|False
a|Filter by statistics.cloud.iops_raw.read
```

* Introduced in: 9.7

```
|statistics.cloud.iops_raw.other
|integer
|query
|False
a|Filter by statistics.cloud.iops_raw.other
```

* Introduced in: 9.7

```
|statistics.cloud.iops_raw.write
|integer
|query
|False
a|Filter by statistics.cloud.iops_raw.write
```

* Introduced in: 9.7

```
|statistics.cloud.iops_raw.total
|integer
|query
|False
a|Filter by statistics.cloud.iops_raw.total
```

* Introduced in: 9.7

```
|statistics.cloud.latency_raw.read
|integer
|query
|False
a|Filter by statistics.cloud.latency_raw.read
```

* Introduced in: 9.7

```
|statistics.cloud.latency_raw.other
|integer
|query
|False
a|Filter by statistics.cloud.latency_raw.other
```

* Introduced in: 9.7

```
|statistics.cloud.latency_raw.write
|integer
|query
|False
a|Filter by statistics.cloud.latency_raw.write
```

* Introduced in: 9.7

```
|statistics.cloud.latency_raw.total
|integer
|query
|False
a|Filter by statistics.cloud.latency_raw.total

* Introduced in: 9.7

|statistics.nfs_ops_raw.link.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.link.total_time

* Introduced in: 9.11

|statistics.nfs_ops_raw.link.count
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.link.count

* Introduced in: 9.11

|statistics.nfs_ops_raw.unlink.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.unlink.total_time

* Introduced in: 9.11

|statistics.nfs_ops_raw.unlink.count
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.unlink.count

* Introduced in: 9.11

|statistics.nfs_ops_raw.lookup.total_time
|integer
```

```
|query
|False
a|Filter by statistics.nfs_ops_raw.lookup.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.lookup.count
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.lookup.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.readlink.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.readlink.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.readlink.count
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.readlink.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.open.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.open.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.open.count
|integer
|query
|False
```

```
a|Filter by statistics.nfs_ops_raw.open.count

* Introduced in: 9.11

|statistics.nfs_ops_raw.watch.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.watch.total_time

* Introduced in: 9.11

|statistics.nfs_ops_raw.watch.count
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.watch.count

* Introduced in: 9.11

|statistics.nfs_ops_raw.readdir.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.readdir.total_time

* Introduced in: 9.11

|statistics.nfs_ops_raw.readdir.count
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.readdir.count

* Introduced in: 9.11

|statistics.nfs_ops_raw.rename.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.rename.total_time
```


* Introduced in: 9.11

```
|statistics.nfs_ops_raw.rename.count  
|integer  
|query  
|False  
a|Filter by statistics.nfs_ops_raw.rename.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.create.symlink.total_time  
|integer  
|query  
|False  
a|Filter by statistics.nfs_ops_raw.create.symlink.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.create.symlink.count  
|integer  
|query  
|False  
a|Filter by statistics.nfs_ops_raw.create.symlink.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.create.dir.total_time  
|integer  
|query  
|False  
a|Filter by statistics.nfs_ops_raw.create.dir.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.create.dir.count  
|integer  
|query  
|False  
a|Filter by statistics.nfs_ops_raw.create.dir.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.create.file.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.create.file.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.create.file.count
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.create.file.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.create.other.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.create.other.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.create.other.count
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.create.other.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.audit.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.audit.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.audit.count
```

```
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.audit.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.access.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.access.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.access.count
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.access.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.write.count
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.write.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.write.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.write.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.write.volume_protocol_latency_histogram_labels
|string
|query
```

```
|False  
a|Filter by  
statistics.nfs_ops_raw.write.volume_protocol_latency_histogram_labels
```

```
* Introduced in: 9.11
```

```
|statistics.nfs_ops_raw.write.volume_protocol_size_histogram_labels  
|string  
|query  
|False  
a|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_counts  
|integer  
|query  
|False  
a|Filter by  
statistics.nfs_ops_raw.write.volume_protocol_latency_histogram_counts
```

```
* Introduced in: 9.11
```

```
|statistics.nfs_ops_raw.write.volume_protocol_size_histogram_counts  
|integer  
|query  
|False  
a|Filter by  
statistics.nfs_ops_raw.write.volume_protocol_size_histogram_counts
```

```
* Introduced in: 9.11
```

```
|statistics.nfs_ops_raw.getattr.total_time  
|integer  
|query  
|False  
a|Filter by statistics.nfs_ops_raw.getattr.total_time
```

```
* Introduced in: 9.11
```

```
|statistics.nfs_ops_raw.getattr.count
```

```
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.getattr.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.read.count
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.read.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.read.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.read.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.read.volume_protocol_latency_histogram_labels
|string
|query
|False
a|Filter by
statistics.nfs_ops_raw.read.volume_protocol_latency_histogram_labels
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.read.volume_protocol_size_histogram_labels
|string
|query
|False
a|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_counts
```

```
|integer
|query
|False
a|Filter by
statistics.nfs_ops_raw.read.volume_protocol_latency_histogram_counts
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.read.volume_protocol_size_histogram_counts
|integer
|query
|False
a|Filter by
statistics.nfs_ops_raw.read.volume_protocol_size_histogram_counts
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.setattr.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.setattr.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.setattr.count
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.setattr.count
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.lock.total_time
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.lock.total_time
```

* Introduced in: 9.11

```
|statistics.nfs_ops_raw.lock.count
```

```
|integer
|query
|False
a|Filter by statistics.nfs_ops_raw.lock.count
```

* Introduced in: 9.11

```
|statistics.throughput_raw.read
|integer
|query
|False
a|Filter by statistics.throughput_raw.read
```

```
|statistics.throughput_raw.other
|integer
|query
|False
a|Filter by statistics.throughput_raw.other
```

```
|statistics.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.throughput_raw.write
```

```
|statistics.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.throughput_raw.total
```

```
|statistics.flexcache_raw.cache_miss_blocks
|integer
|query
|False
a|Filter by statistics.flexcache_raw.cache_miss_blocks
```

* Introduced in: 9.8

```
|statistics.flexcache_raw.timestamp
|string
```

```
|query
|False
a|Filter by statistics.flexcache_raw.timestamp

* Introduced in: 9.8

|statistics.flexcache_raw.status
|string
|query
|False
a|Filter by statistics.flexcache_raw.status

* Introduced in: 9.8

|statistics.flexcache_raw.client_requested_blocks
|integer
|query
|False
a|Filter by statistics.flexcache_raw.client_requested_blocks

* Introduced in: 9.8

|statistics.iops_raw.read
|integer
|query
|False
a|Filter by statistics.iops_raw.read

|statistics.iops_raw.other
|integer
|query
|False
a|Filter by statistics.iops_raw.other

|statistics.iops_raw.write
|integer
|query
|False
a|Filter by statistics.iops_raw.write

|statistics.iops_raw.total
```



```
|integer
|query
|False
a|Filter by statistics.iops_raw.total

|statistics.latency_raw.read
|integer
|query
|False
a|Filter by statistics.latency_raw.read

|statistics.latency_raw.other
|integer
|query
|False
a|Filter by statistics.latency_raw.other

|statistics.latency_raw.write
|integer
|query
|False
a|Filter by statistics.latency_raw.write

|statistics.latency_raw.total
|integer
|query
|False
a|Filter by statistics.latency_raw.total

|anti_ransomware_state
|string
|query
|False
a|Filter by anti_ransomware_state

* Introduced in: 9.10

|cloud_write_enabled
|boolean
|query
|False
```

a|Filter by cloud_write_enabled

* Introduced in: 9.13

|snapshot_locking_enabled

|boolean

|query

|False

a|Filter by snapshot_locking_enabled

* Introduced in: 9.12

|is_object_store

|boolean

|query

|False

a|Filter by is_object_store

* Introduced in: 9.8

|snapmirror.is_protected

|boolean

|query

|False

a|Filter by snapmirror.is_protected

* Introduced in: 9.7

|snapmirror.destinations.is_ontap

|boolean

|query

|False

a|Filter by snapmirror.destinations.is_ontap

* Introduced in: 9.9

|snapmirror.destinations.is_cloud

|boolean

|query

|False

a|Filter by snapmirror.destinations.is_cloud

* Introduced in: 9.9

|qos.policy.name
|string
|query
|False
a|Filter by qos.policy.name

|qos.policy.min_throughput_mbps
|integer
|query
|False
a|Filter by qos.policy.min_throughput_mbps

* Introduced in: 9.8

|qos.policy.min_throughput_iops
|integer
|query
|False
a|Filter by qos.policy.min_throughput_iops

|qos.policy.max_throughput_mbps
|integer
|query
|False
a|Filter by qos.policy.max_throughput_mbps

|qos.policy.uuid
|string
|query
|False
a|Filter by qos.policy.uuid

|qos.policy.max_throughput_iops
|integer
|query
|False
a|Filter by qos.policy.max_throughput_iops

```
|convert_unicode
|boolean
|query
|False
a|Filter by convert_unicode
```

* Introduced in: 9.10

```
|rebalancing.min_threshold
|integer
|query
|False
a|Filter by rebalancing.min_threshold
```

* Introduced in: 9.11

```
|rebalancing.imbalance_percent
|integer
|query
|False
a|Filter by rebalancing.imbalance_percent
```

* Introduced in: 9.11

```
|rebalancing.runtime
|string
|query
|False
a|Filter by rebalancing.runtime
```

* Introduced in: 9.11

```
|rebalancing.max_constituent_imbalance_percent
|integer
|query
|False
a|Filter by rebalancing.max_constituent_imbalance_percent
```

* Introduced in: 9.11

```
|rebalancing.target_used
|integer
```

```
|query
|False
a|Filter by rebalancing.target_used

* Introduced in: 9.11

|rebalancing.min_file_size
|integer
|query
|False
a|Filter by rebalancing.min_file_size

* Introduced in: 9.11

|rebalancing.imbalance_size
|integer
|query
|False
a|Filter by rebalancing.imbalance_size

* Introduced in: 9.11

|rebalancing.notices.arguments.message
|string
|query
|False
a|Filter by rebalancing.notices.arguments.message

* Introduced in: 9.12

|rebalancing.notices.arguments.code
|string
|query
|False
a|Filter by rebalancing.notices.arguments.code

* Introduced in: 9.12

|rebalancing.notices.message
|string
|query
|False
```

a|Filter by rebalancing.notices.message

* Introduced in: 9.12

|rebalancing.notices.code

|string

|query

|False

a|Filter by rebalancing.notices.code

* Introduced in: 9.12

|rebalancing.notices.target

|string

|query

|False

a|Filter by rebalancing.notices.target

* Introduced in: 9.12

|rebalancing.stop_time

|string

|query

|False

a|Filter by rebalancing.stop_time

* Introduced in: 9.11

|rebalancing.max_threshold

|integer

|query

|False

a|Filter by rebalancing.max_threshold

* Introduced in: 9.11

|rebalancing.max_file_moves

|integer

|query

|False

a|Filter by rebalancing.max_file_moves

* Introduced in: 9.11

```
|rebalancing.engine.scanner.blocks_scanned  
|integer  
|query  
|False  
a|Filter by rebalancing.engine.scanner.blocks_scanned
```

* Introduced in: 9.12

```
|rebalancing.engine.scanner.blocks_skipped.on_demand_destination  
|integer  
|query  
|False  
a|Filter by  
rebalancing.engine.scanner.blocks_skipped.on_demand_destination
```

* Introduced in: 9.12

```
|rebalancing.engine.scanner.blocks_skipped.efficiency_blocks  
|integer  
|query  
|False  
a|Filter by rebalancing.engine.scanner.blocks_skipped.efficiency_blocks
```

* Introduced in: 9.12

```
|rebalancing.engine.scanner.blocks_skipped.other  
|integer  
|query  
|False  
a|Filter by rebalancing.engine.scanner.blocks_skipped.other
```

* Introduced in: 9.12

```
|rebalancing.engine.scanner.blocks_skipped.too_small  
|integer  
|query  
|False  
a|Filter by rebalancing.engine.scanner.blocks_skipped.too_small
```

* Introduced in: 9.12

```
|rebalancing.engine.scanner.blocks_skipped.metadata
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.blocks_skipped.metadata
```

* Introduced in: 9.12

```
|rebalancing.engine.scanner.blocks_skipped.write_fenced
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.blocks_skipped.write_fenced
```

* Introduced in: 9.12

```
|rebalancing.engine.scanner.blocks_skipped.too_large
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.blocks_skipped.too_large
```

* Introduced in: 9.12

```
|rebalancing.engine.scanner.blocks_skipped.in_snapshot
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.blocks_skipped.in_snapshot
```

* Introduced in: 9.12

```
|rebalancing.engine.scanner.blocks_skipped.fast_truncate
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.blocks_skipped.fast_truncate
```

* Introduced in: 9.12


```
|rebalancing.engine.scanner.blocks_skipped.footprint_invalid
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.blocks_skipped.footprint_invalid

* Introduced in: 9.12

|rebalancing.engine.scanner.blocks_skipped.efficiency_percent
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.blocks_skipped.efficiency_percent

* Introduced in: 9.12

|rebalancing.engine.scanner.blocks_skipped.incompatible
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.blocks_skipped.incompatible

* Introduced in: 9.12

|rebalancing.engine.scanner.blocks_skipped.remote_cache
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.blocks_skipped.remote_cache

* Introduced in: 9.12

|rebalancing.engine.scanner.files_skipped.too_large
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.files_skipped.too_large

* Introduced in: 9.12

|rebalancing.engine.scanner.files_skipped.write_fenced
|integer
```

```
|query
|False
a|Filter by rebalancing.engine.scanner.files_skipped.write_fenced

* Introduced in: 9.12

|rebalancing.engine.scanner.files_skipped.metadata
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.files_skipped.metadata

* Introduced in: 9.12

|rebalancing.engine.scanner.files_skipped.too_small
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.files_skipped.too_small

* Introduced in: 9.12

|rebalancing.engine.scanner.files_skipped.other
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.files_skipped.other

* Introduced in: 9.12

|rebalancing.engine.scanner.files_skipped.efficiency_blocks
|integer
|query
|False
a|Filter by rebalancing.engine.scanner.files_skipped.efficiency_blocks

* Introduced in: 9.12

|rebalancing.engine.scanner.files_skipped.on_demand_destination
|integer
|query
|False
```

a|Filter by rebalancing.engine.scanner.files_skipped.on_demand_destination

* Introduced in: 9.12

|rebalancing.engine.scanner.files_skipped.remote_cache

|integer

|query

|False

a|Filter by rebalancing.engine.scanner.files_skipped.remote_cache

* Introduced in: 9.12

|rebalancing.engine.scanner.files_skipped.incompatible

|integer

|query

|False

a|Filter by rebalancing.engine.scanner.files_skipped.incompatible

* Introduced in: 9.12

|rebalancing.engine.scanner.files_skipped.efficiency_percent

|integer

|query

|False

a|Filter by rebalancing.engine.scanner.files_skipped.efficiency_percent

* Introduced in: 9.12

|rebalancing.engine.scanner.files_skipped.footprint_invalid

|integer

|query

|False

a|Filter by rebalancing.engine.scanner.files_skipped.footprint_invalid

* Introduced in: 9.12

|rebalancing.engine.scanner.files_skipped.in_snapshot

|integer

|query

|False

a|Filter by rebalancing.engine.scanner.files_skipped.in_snapshot

* Introduced in: 9.12

```
|rebalancing.engine.scanner.files_skipped.fast_truncate  
|integer  
|query  
|False  
a|Filter by rebalancing.engine.scanner.files_skipped.fast_truncate
```

* Introduced in: 9.12

```
|rebalancing.engine.scanner.files_scanned  
|integer  
|query  
|False  
a|Filter by rebalancing.engine.scanner.files_scanned
```

* Introduced in: 9.12

```
|rebalancing.engine.movement.last_error.time  
|string  
|query  
|False  
a|Filter by rebalancing.engine.movement.last_error.time
```

* Introduced in: 9.12

```
|rebalancing.engine.movement.last_error.file_id  
|integer  
|query  
|False  
a|Filter by rebalancing.engine.movement.last_error.file_id
```

* Introduced in: 9.12

```
|rebalancing.engine.movement.last_error.code  
|integer  
|query  
|False  
a|Filter by rebalancing.engine.movement.last_error.code
```

* Introduced in: 9.12

```
|rebalancing.engine.movement.last_error.destination
|integer
|query
|False
a|Filter by rebalancing.engine.movement.last_error.destination
```

* Introduced in: 9.12

```
|rebalancing.engine.movement.file_moves_started
|integer
|query
|False
a|Filter by rebalancing.engine.movement.file_moves_started
```

* Introduced in: 9.12

```
|rebalancing.engine.movement.most_recent_start_time
|string
|query
|False
a|Filter by rebalancing.engine.movement.most_recent_start_time
```

* Introduced in: 9.12

```
|rebalancing.state
|string
|query
|False
a|Filter by rebalancing.state
```

* Introduced in: 9.11

```
|rebalancing.used_for_imbalance
|integer
|query
|False
a|Filter by rebalancing.used_for_imbalance
```

* Introduced in: 9.12

```
|rebalancing.data_moved
```

```
|integer
|query
|False
a|Filter by rebalancing.data_moved

* Introduced in: 9.11

|rebalancing.exclude_snapshots
|boolean
|query
|False
a|Filter by rebalancing.exclude_snapshots

* Introduced in: 9.11

|rebalancing.start_time
|string
|query
|False
a|Filter by rebalancing.start_time

* Introduced in: 9.11

|rebalancing.max_runtime
|string
|query
|False
a|Filter by rebalancing.max_runtime

* Introduced in: 9.11

|movement.percent_complete
|integer
|query
|False
a|Filter by movement.percent_complete

|movement.cutover_window
|integer
|query
|False
a|Filter by movement.cutover_window
```

```
|movement.destination_aggregate.name
|string
|query
|False
a|Filter by movement.destination_aggregate.name
```

```
|movement.destination_aggregate.uuid
|string
|query
|False
a|Filter by movement.destination_aggregate.uuid
```

```
|movement.state
|string
|query
|False
a|Filter by movement.state
```

```
|movement.start_time
|string
|query
|False
a|Filter by movement.start_time
```

* Introduced in: 9.9

```
|max_dir_size
|integer
|query
|False
a|Filter by max_dir_size
```

* Introduced in: 9.10

```
|analytics.supported
|boolean
|query
|False
a|Filter by analytics.supported
```

* Introduced in: 9.8

|analytics.state
|string
|query
|False
a|Filter by analytics.state

* Introduced in: 9.8

|analytics.unsupported_reason.message
|string
|query
|False
a|Filter by analytics.unsupported_reason.message

* Introduced in: 9.8

|analytics.unsupported_reason.code
|string
|query
|False
a|Filter by analytics.unsupported_reason.code

* Introduced in: 9.8

|analytics.scan_progress
|integer
|query
|False
a|Filter by analytics.scan_progress

* Introduced in: 9.8

|analytics.initialization.state
|string
|query
|False
a|Filter by analytics.initialization.state

* Introduced in: 9.12


```
|status
|string
|query
|False
a|Filter by status
```

* Introduced in: 9.9

```
|constituents.name
|string
|query
|False
a|Filter by constituents.name
```

* Introduced in: 9.9

```
|constituents.space.available_percent
|integer
|query
|False
a|Filter by constituents.space.available_percent
```

* Introduced in: 9.9

```
|constituents.space.used_percent
|integer
|query
|False
a|Filter by constituents.space.used_percent
```

* Introduced in: 9.10

```
|constituents.space.afs_total
|integer
|query
|False
a|Filter by constituents.space.afs_total
```

* Introduced in: 9.9

```
|constituents.space.local_tier_footprint
```

```
|integer
|query
|False
a|Filter by constituents.space.local_tier_footprint
```

* Introduced in: 9.9

```
|constituents.space.used_by_afs
|integer
|query
|False
a|Filter by constituents.space.used_by_afs
```

* Introduced in: 9.9

```
|constituents.space.large_size_enabled
|boolean
|query
|False
a|Filter by constituents.space.large_size_enabled
```

* Introduced in: 9.12

```
|constituents.space.used
|integer
|query
|False
a|Filter by constituents.space.used
```

* Introduced in: 9.9

```
|constituents.space.snapshot.used
|integer
|query
|False
a|Filter by constituents.space.snapshot.used
```

* Introduced in: 9.9

```
|constituents.space.snapshot.reserve_percent
|integer
|query
```

```
|False
a|Filter by constituents.space.snapshot.reserve_percent

* Introduced in: 9.9

|constituents.space.performance_tier_footprint
|integer
|query
|False
a|Filter by constituents.space.performance_tier_footprint

* Introduced in: 9.9

|constituents.space.footprint
|integer
|query
|False
a|Filter by constituents.space.footprint

* Introduced in: 9.9

|constituents.space.capacity_tier_footprint
|integer
|query
|False
a|Filter by constituents.space.capacity_tier_footprint

* Introduced in: 9.9

|constituents.space.over_provisioned
|integer
|query
|False
a|Filter by constituents.space.over_provisioned

* Introduced in: 9.9

|constituents.space.logical_space.available
|integer
|query
|False
a|Filter by constituents.space.logical_space.available
```

* Introduced in: 9.9

|constituents.space.logical_space.enforcement

|boolean

|query

|False

a|Filter by constituents.space.logical_space.enforcement

* Introduced in: 9.9

|constituents.space.logical_space.reporting

|boolean

|query

|False

a|Filter by constituents.space.logical_space.reporting

* Introduced in: 9.9

|constituents.space.logical_space.used_by_afs

|integer

|query

|False

a|Filter by constituents.space.logical_space.used_by_afs

* Introduced in: 9.9

|constituents.space.metadata

|integer

|query

|False

a|Filter by constituents.space.metadata

* Introduced in: 9.9

|constituents.space.block_storage_inactive_user_data

|integer

|query

|False

a|Filter by constituents.space.block_storage_inactive_user_data

* Introduced in: 9.9

```
|constituents.space.size
|integer
|query
|False
a|Filter by constituents.space.size
```

* Introduced in: 9.9

```
|constituents.space.total_footprint
|integer
|query
|False
a|Filter by constituents.space.total_footprint
```

* Introduced in: 9.9

```
|constituents.space.available
|integer
|query
|False
a|Filter by constituents.space.available
```

* Introduced in: 9.9

```
|constituents.aggregates.name
|string
|query
|False
a|Filter by constituents.aggregates.name
```

* Introduced in: 9.9

```
|constituents.aggregates.uuid
|string
|query
|False
a|Filter by constituents.aggregates.uuid
```

* Introduced in: 9.9

```
|constituents.movement.percent_complete
|integer
|query
|False
a|Filter by constituents.movement.percent_complete

* Introduced in: 9.9

|constituents.movement.cutover_window
|integer
|query
|False
a|Filter by constituents.movement.cutover_window

* Introduced in: 9.9

|constituents.movement.state
|string
|query
|False
a|Filter by constituents.movement.state

* Introduced in: 9.9

|constituents.movement.destination_aggregate.name
|string
|query
|False
a|Filter by constituents.movement.destination_aggregate.name

* Introduced in: 9.9

|constituents.movement.destination_aggregate.uuid
|string
|query
|False
a|Filter by constituents.movement.destination_aggregate.uuid

* Introduced in: 9.9

|snapshot_count
|integer
```

```
|query
|False
a|Filter by snapshot_count

* Introduced in: 9.10
* Max value: 1023
* Min value: 0

|files.maximum
|integer
|query
|False
a|Filter by files.maximum

|files.used
|integer
|query
|False
a|Filter by files.used

|style
|string
|query
|False
a|Filter by style

|flexcache_endpoint_type
|string
|query
|False
a|Filter by flexcache_endpoint_type

|msid
|integer
|query
|False
a|Filter by msid

* Introduced in: 9.11

|tiering.min_cooling_days
```

```
|integer
|query
|False
a|Filter by tiering.min_cooling_days
```

```
* Introduced in: 9.8
* Max value: 183
* Min value: 2
```

```
|tiering.policy
|string
|query
|False
a|Filter by tiering.policy
```

```
|tiering.object_tags
|string
|query
|False
a|Filter by tiering.object_tags
```

```
* Introduced in: 9.8
* maxLength: 257
```

```
|name
|string
|query
|False
a|Filter by name
```

```
* maxLength: 203
* minLength: 1
```

```
|state
|string
|query
|False
a|Filter by state
```

```
|uuid
|string
|query
```



```
|False  
a|Filter by uuid
```

```
|fields  
|array[string]  
|query  
|False  
a|Specify the fields to return.
```

```
|max_records  
|integer  
|query  
|False  
a|Limit the number of records returned.
```

```
|return_records  
|boolean  
|query  
|False  
a|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  
a|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  
a|Order results by specified fields and optional [asc|desc] direction.
```

Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

```
[cols=3*,options=header]
```

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|num_records

|integer

a|Number of records

|records

|array[link:#volume[volume]]

a|

|===

.Example response

[%collapsible%closed]

====

```
[source,json,subs=+macros]
```

```
{
```

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

```
    "time": "2021-06-01 15:06:41 +0000"
  },
  "dry_run_start_time": "string",
  "space": {
    "snapshot_count": 0,
    "used": 0,
    "used_by_logs": 0,
    "used_by_snapshots": 0
  },
  "state": "disabled",
  "suspect_files": {
    "count": 0,
    "entropy": "string",
    "format": "string"
  }
},
"anti_ransomware_state": "disabled",
"application": {
  "name": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
},
"autosize": {
  "mode": "grow"
},
"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": "default",
"comment": "string",
"consistency_group": {
  "name": "consistency_group_1",
  "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
},
"constituents": {
  "aggregates": {
    "name": "string",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  },
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": 0,
    "state": "replicating",
    "tiering_policy": "all"
  },
  "name": "string",
  "space": {
    "available": 0,
    "block_storage_inactive_user_data": 0,
    "capacity_tier_footprint": 0,
    "footprint": 0,
    "local_tier_footprint": 0,
    "logical_space": {
      "available": 0,

```

```

    "used_by_afs": 0
  },
  "metadata": 0,
  "over_provisioned": 0,
  "performance_tier_footprint": 0,
  "snapshot": {
    "used": 0
  },
  "total_footprint": 0,
  "used": 0
}
},
"create_time": "2018-06-04 19:00:00 +0000",
"efficiency": {
  "application_io_size": "8k",
  "auto_state": "auto",
  "compaction": "inline",
  "compression": "inline",
  "compression_type": "none",
  "cross_volume_dedupe": "inline",
  "dedupe": "inline",
  "idcs_scanner": {
    "mode": "default",
    "operation_state": "idle",
    "status": "success",
    "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": "idle",
  "progress": "string",
  "scanner": {
    "state": "idle"
  },
  "schedule": "string",
  "space_savings": {
    "compression": 0,
    "compression_percent": 0,
    "dedupe": 0,
    "dedupe_percent": 0,
    "dedupe_sharing": 0,
    "total": 0,
    "total_percent": 0
  }
}

```

```

    },
    "state": "disabled",
    "storage_efficiency_mode": "default",
    "type": "regular",
    "volume_path": "string"
  },
  "encryption": {
    "key_create_time": "2022-01-01 19:00:00 +0000",
    "key_id": "string",
    "key_manager_attribute": "CRN=v1:bluemix:public:containers-
kubernetes:us-south:a/asdfghjkl1234:asdfghjkl1234:worker:kubernetes-
asdfghjkl-worker1",
    "state": "encrypted",
    "status": {
      "code": "string",
      "message": "string"
    },
    "type": "none"
  },
  "files": {
    "used": 0
  },
  "flash_pool": {
    "cache_eligibility": "read",
    "cache_retention_priority": "normal",
    "caching_policy": "none"
  },
  "flexcache_endpoint_type": "none",
  "flexgroup": {
    "name": "my_flexgroup",
    "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
  },
  "guarantee": {
    "type": "volume"
  },
  "language": "ar",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "cloud": {
    "duration": "PT15S",
    "iops": {
      "read": 200,

```

```

    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"duration": "PT15S",
"flexcache": {
  "bandwidth_savings": 4096,
  "cache_miss_percent": 20,
  "duration": "PT1D",
  "status": "ok",
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"iops": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"latency": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"status": "ok",
"throughput": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 +0000"
},
"movement": {
  "cutover_window": 30,
  "destination_aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "aggr1",

```



```

    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "percent_complete": 0,
  "start_time": "2020-12-07 08:45:12 +0000",
  "state": "replicating",
  "tiering_policy": "all"
},
"name": "vol_cs_dept",
"nas": {
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "junction_parent": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "vs1_root",
    "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
  },
  "path": "/user/my_volume",
  "security_style": "mixed",
  "unix_permissions": 493
},
"qos": {
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},

```

```
"quota": {
  "state": "corrupt"
},
"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 19:00:00 +0000"
      },
      "most_recent_start_time": "2018-06-04 19:00:00 +0000"
    },
    "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,
"notices": {
    "arguments": {
        "code": "string",
        "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
},
"runtime": "string",
"state": "rebalancing",
"stop_time": "string",
"target_used": 0,
"used_for_imbalance": 0
},
"scheduled_snapshot_naming_scheme": "create_time",
"snaplock": {
    "append_mode_enabled": "",
    "autocommit_period": "P30M",
    "compliance_clock_time": "2018-06-04 19:00:00 +0000",
    "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,
  "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
  },
  "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": "try",
      "defer_delete": "scheduled",
      "delete_order": "newest_first",

```

```

    "prefix": "string",
    "trigger": "volume"
  },
  "autodelete_trigger": "volume",
  "reserve_available": 0,
  "reserve_size": 0,
  "space_used_percent": 0,
  "used": 0
},
"snapshot_reserve_unusable": 0,
"snapshot_spill": 0,
"total_footprint": 0,
"used": 0,
"user_data": 0,
"volume_guarantee_footprint": 0
},
"state": "error",
"statistics": {
  "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,
```



```
"<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,
    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 11:20:13 +0000"
},
"flexcache_raw": {
    "cache_miss_blocks": 10,
    "client_requested_blocks": 500,
    "status": "ok",
    "timestamp": "2017-01-25 11:20:13 +0000"
},
"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,
```



```
"<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
}

```



```
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 11:20:13 +0000"
},
"status": {
},
```

```

"style": "flexvol",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"tiering": {
  "object_tags": {
  },
  "policy": "all",
  "storage_class": "default"
},
"type": "rw",
"uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {

```

```

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

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|next

```

```

|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#unsupported_reason]
[.api-collapsible-fifth-title]
unsupported_reason

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|If volume activity tracking is not supported on the volume, this field
provides an appropriate error code.

|message
|string
a|If volume activity tracking is not supported on the volume, this field
provides an error message detailing why this is the case.

```

```
|===
```

```
[#activity_tracking]  
[.api-collapsible-fifth-title]  
activity_tracking
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|state
```

```
|string
```

```
a|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
```

```
a|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
```

```
|link:#unsupported_reason[unsupported_reason]
```

```
a|
```

```
|===
```

```
[#aggregates]  
[.api-collapsible-fifth-title]  
aggregates
```

Aggregate

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#initialization]
```

```
[.api-collapsible-fifth-title]
```

```
initialization
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|state
```

```
|string
```

```
a|State of the analytics file system scan.
```

```
|===
```

```
[#unsupported_reason]
```

```
[.api-collapsible-fifth-title]
```

```
unsupported_reason
```

```
[cols=3*,options=header]
```



```

|===
|Name
|Type
|Description

|code
|string
a|If file system analytics is not supported on the volume, this field
provides the error code explaining why.

|message
|string
a|If file system analytics is not supported on the volume, this field
provides the error message explaining why.

|===

[#analytics]
[.api-collapsible-fifth-title]
analytics

[cols=3*,options=header]
|===
|Name
|Type
|Description

|initialization
|link:#initialization[initialization]
a|

|scan_progress
|integer
a|Percentage of files in the volume that the file system analytics
initialization scan has processed. Only returned when the state is
`initializing`.

|state
|string
a|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

a|This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field.

|unsupported_reason

|link:#unsupported_reason[unsupported_reason]

a|

|===

[#_links]

[.api-collapsible-fifth-title]

_links

[cols=3*,options=header]

|===

|Name

|Type

|Description

|suspects

|link:#href[href]

a|

|===

[#anti_ransomware_attack_report]

[.api-collapsible-fifth-title]

```
anti_ransomware_attack_report
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|time
```

```
|string
```

```
a|Timestamp at which ransomware attack is observed.
```

```
|===
```

```
[#space]
```

```
[.api-collapsible-fifth-title]
```

```
space
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|snapshot_count
```

```
|integer
```

```
a|Total number of Anti-ransomware backup Snapshot copies.
```

```
|used
```

```
|integer
```

```
a|Total space in bytes used by the Anti-ransomware feature.
```

```
|used_by_logs
```

```
|integer
```

```
a|Space in bytes used by the Anti-ransomware analytics logs.
```

```
|used_by_snapshots
```

```
|integer
```

a|Space in bytes used by the Anti-ransomware backup Snapshot copies.

|===

[#suspect_files]

[.api-collapsible-fifth-title]

suspect_files

[cols=3*,options=header]

|===

|Name

|Type

|Description

|count

|integer

a|Total number of `suspect_files.format` files observed by the Anti-ransomware analytics engine on the volume.

|entropy

|string

a|Indicates the entropy level of this file type.

|format

|string

a|File formats observed by the Anti-ransomware analytics engine on the volume.

|===

[#anti_ransomware]

[.api-collapsible-fifth-title]

anti_ransomware

Anti-ransomware related information of the volume.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|attack_probability

|string

a|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[link:#anti_ransomware_attack_report[anti_ransomware_attack_report]]

a|

|dry_run_start_time

|string

a|Time when Anti-ransomware monitoring `state` is set to dry-run value for starting evaluation mode.

|space

|link:#space[space]

a|

|state

|string

a|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
a|Indicates whether or not to set the surge values as historical values.
```

```
|suspect_files
|array[link:#suspect_files[suspect_files]]
a|
```

```
|===
```

```
[#application]
[.api-collapsible-fifth-title]
application
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|Name of the application to which the volume belongs. Available only when
the volume is part of an application.
```

```
|uuid
|string
a|UUID of the application to which the volume belongs. Available only when
the volume is part of an application.
```

```
|===
```

```
[#asynchronous_directory_delete]
[.api-collapsible-fifth-title]
asynchronous_directory_delete
```

Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

|Type

|Description

|enabled

|boolean

a|Specifies whether asynchronous directory delete from the client is enabled on the volume.

|trash_bin

|string

a|Name of the trash bin directory. If no "trash_bin" property is specified when enabling, the default trash bin name, "._ontaptrashbin", is used.

|===

[#autosize]

[.api-collapsible-fifth-title]

autosize

[cols=3*,options=header]

|===

|Name

|Type

|Description

|grow_threshold

|integer

a|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

a|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

a|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

a|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.

|shrink_threshold

|integer

a|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]

[.api-collapsible-fifth-title]

snapshot_reference

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|

|uuid

|string


```

a|

|===

[#parent_svm]
[.api-collapsible-fifth-title]
parent_svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#parent_volume]
[.api-collapsible-fifth-title]
parent_volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name

```

```

|string
a|The name of the volume.

|uuid
|string
a|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]
[.api-collapsible-fifth-title]
clone

[cols=3*,options=header]
|===
|Name
|Type
|Description

|inherited_physical_used
|integer
a|Inherited physical used from the clone's base snapshot.

|inherited_savings
|integer
a|Inherited savings from the clone's base snapshot.

|is_flexclone
|boolean
a|Specifies if this volume is a normal FlexVol or FlexClone. This field
needs to be set when creating a FlexClone. Valid in POST.

|parent_snapshot
|link:#snapshot_reference[snapshot_reference]
a|

```

```
|parent_svm
|link:#parent_svm[parent_svm]
a|
```

```
|parent_volume
|link:#parent_volume[parent_volume]
a|
```

```
|split_complete_percent
|integer
a|Percentage of FlexClone blocks split from its parent volume.
```

```
|split_estimate
|integer
a|Space required by the containing-aggregate to split the FlexClone volume.
```

```
|split_initiated
|boolean
a|This field is set when split is executed on any FlexClone, that is when the FlexClone volume is split from its parent FlexVol. This field needs to be set for splitting a FlexClone form FlexVol. Valid in PATCH.
```

```
|===
```

```
[#consistency_group]
[.api-collapsible-fifth-title]
consistency_group
```

Consistency group the volume is part of.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
```

```
|string
```

```
a|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

a|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]

[.api-collapsible-fifth-title]

aggregates

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|Name of the aggregate hosting the FlexGroup Constituent.

|uuid

|string

a|Unique identifier for the aggregate.

|===

[#destination_aggregate]

[.api-collapsible-fifth-title]

destination_aggregate

Aggregate

[cols=3*,options=header]

|===

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#movement]
[.api-collapsible-fifth-title]
movement
```

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|cutover_window
|integer
```

a|Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.

```
|destination_aggregate
|link:#destination_aggregate[destination_aggregate]
a|Aggregate
```

```
|percent_complete
|integer
```

a|Completion percentage

|state

|string

a|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

a|Tiering policy for FabricPool

|===

[#logical_space]

[.api-collapsible-fifth-title]

logical_space

[cols=3*,options=header]

|===

|Name

|Type

|Description

|available

|integer

a|The amount of space available in this volume with storage efficiency space considered used, in bytes.

|enforcement

|boolean

a|Specifies whether space accounting for operations on the volume is done along with storage efficiency.

|reporting

|boolean

a|Specifies whether space reporting on the volume is done along with storage efficiency.

|used_by_afs

|integer

a|The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes.

|===

[#snapshot]

[.api-collapsible-fifth-title]

snapshot

[cols=3*,options=header]

|===

|Name

|Type

|Description

|autodelete_enabled

|boolean

a|Specifies whether Snapshot copy autodelete is currently enabled on this volume.

|reserve_percent

|integer

a|The space that has been set aside as a reserve for Snapshot copy usage, in percent.

|used

|integer

a|The total space used by Snapshot copies in the volume, in bytes.

|===

[#space]

[.api-collapsible-fifth-title]

space

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|afs_total
```

```
|integer
```

```
a|Total size of AFS, excluding snap-reserve, in bytes.
```

```
|available
```

```
|integer
```

```
a|The available space, in bytes.
```

```
|available_percent
```

```
|integer
```

```
a|The space available, as a percent.
```

```
|block_storage_inactive_user_data
```

```
|integer
```

```
a|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
```

```
a|Space used by capacity tier for this volume in the FabricPool aggregate, in bytes.
```

```
|footprint
```

```
|integer
```

```
a|Data used for this volume in the aggregate, in bytes.
```

```
|large_size_enabled
```

```
|boolean
```

```
a|Specifies whether the support for large volumes and large files is enabled on the volume.
```

```
|local_tier_footprint
```



```
|integer
a|Space used by the local tier for this volume in the aggregate, in bytes.

|logical_space
|link:#logical_space[logical_space]
a|

|metadata
|integer
a|Space used by the volume metadata in the aggregate, in bytes.

|over_provisioned
|integer
a|The amount of space not available for this volume in the aggregate, in
bytes.

|performance_tier_footprint
|integer
a|Space used by the performance tier for this volume in the FabricPool
aggregate, in bytes.

|size
|integer
a|Total provisioned size. The default size is equal to the minimum size of
20MB, in bytes.

|snapshot
|link:#snapshot[snapshot]
a|

|total_footprint
|integer
a|Data and metadata used for this volume in the aggregate, in bytes.

|used
|integer
a|The virtual space used (includes volume reserves) before storage
efficiency, in bytes.

|used_by_afs
```

```

|integer
a|The space used by Active Filesystem, in bytes.

|used_percent
|integer
a|The virtual space used (includes volume reserves) before storage
efficiency, as a percent.

|===

[#constituents]
[.api-collapsible-fifth-title]
constituents

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregates
|link:#aggregates[aggregates]
a|

|movement
|link:#movement[movement]
a|Volume movement. All attributes are modify, that is, not writable
through POST. Set PATCH state to destination_aggregate to initiate a
volume move operation. Volume movement on FlexGroup constituents are not
supported.

|name
|string
a|FlexGroup Constituents name.

|space
|link:#space[space]
a|

|===

```

```
[#idcs_scanner]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

a|Specifies the administrative state of the inactive data compression scanner.

```
|inactive_days
```

```
|integer
```

a|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
```

a|Specifies the mode of inactive data compression scanner. Valid for PATCH and GET.

```
|operation_state
```

```
|string
```

a|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
```

a|Status of last inactive data compression scan on the volume.

```
|threshold_inactive_time
|string
a|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.+++</num>+++
```

```
|===
```

```
[#policy]
[.api-collapsible-fifth-title]
policy
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|Specifies the name of the efficiency policy.
```

```
|===
```

```
[#scanner]
[.api-collapsible-fifth-title]
scanner
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|compression
|boolean
a|Start compression if scanning old data. Valid for PATCH and GET. This
option is not supported for FSX/CVO platforms.
```

```
|dedupe
```

```

|boolean
a|Start deduplication if scanning old data. Valid for PATCH and GET.

|scan_old_data
|boolean
a|Indicates whether or not to scan old data. Valid for PATCH and GET.

|state
|string
a|State of the volume efficiency scanner. Valid for PATCH and GET. Valid
options for PATCH are "idle" and "active".

|===

[#space_savings]
[.api-collapsible-fifth-title]
space_savings

[cols=3*,options=header]
|===
|Name
|Type
|Description

|compression
|integer
a|Total disk space that is saved by compressing blocks on the referenced
file system, in bytes.

|compression_percent
|integer
a|Percentage of total disk space that is saved by compressing blocks on
the referenced file system.

|dedupe
|integer
a|Total disk space that is saved by deduplication and file cloning, in
bytes.

|dedupe_percent

```

```

|integer
a|Percentage of total disk space that is saved by deduplication and file
cloning.

|dedupe_sharing
|integer
a|Total disk space that is shared due to deduplication and file cloning.

|total
|integer
a|Total disk space saved in the volume due to deduplication, compression
and file cloning, in bytes.

|total_percent
|integer
a|Percentage of total disk space saved in the volume due to deduplication,
compression and file cloning.

|===

[#efficiency]
[.api-collapsible-fifth-title]
efficiency

[cols=3*,options=header]
|===
|Name
|Type
|Description

|application_io_size
|string
a|Block size to use by compression.

|auto_state
|string
a|Automatic deduplication schedule volume state.
auto &dash; Volumes with auto_state set to auto start post-process
deduplication automatically.
deprioritized &dash; Volumes with auto_state set to deprioritized do not
start post-process deduplication automatically.

```

|compaction
|string
a|The system can be enabled/disabled compaction.
inline ‐ Data will be compacted first and written to the volume.
none ‐ None
mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.

|compression
|string
a|The system can be enabled/disabled compression.
inline ‐ Data will be compressed first and written to the volume.
background ‐ Data will be written to the volume and compressed later.
both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run.
none ‐ None
mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.

NOTE: that On volumes with container compression enabled, background compression refers to inactive data compression scan enabled on the volume.

|compression_type
|string
a|Compression type to use by compression. Valid for PATCH and GET.

|cross_volume_dedupe
|string
a|The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled.
inline ‐ Data will be cross volume deduped first and written to the volume.
background ‐ Data will be written to the volume and cross volume deduped later.
both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run.
none ‐ None
mixed ‐ Read only field for FlexGroups, where some of the constituent

volumes are cross volume dedupe enabled and some are disabled.

|dedupe

|string

a|The system can be enabled/disabled dedupe.

inline ‐ Data will be deduped first and written to the volume.

background ‐ Data will be written to the volume and deduped later.

both ‐ Inline dedupe dedupes the data and write to the volume,

background dedupe dedupes only the blocks on which inline dedupe is not run.

none ‐ None

mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.

|has_savings

|boolean

a|When true, indicates that the volume contains shared(deduplication, file clones) or compressed data.

|idcs_scanner

|link:#idcs_scanner[idcs_scanner]

a|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

a|Last sis operation begin timestamp.

|last_op_end

|string

a|Last sis operation end timestamp.

|last_op_err

|string

a|Last sis operation error text.

|last_op_size


```
|integer
a|Last sis operation size.

|last_op_state
|string
a|Last sis operation state.

|logging_enabled
|boolean
a|When true, indicates that space savings for any newly-written data are
being logged.

|op_state
|string
a|Sis status of the volume.

|policy
|link:#policy[policy]
a|

|progress
|string
a|Sis progress of the volume.

|scanner
|link:#scanner[scanner]
a|

|schedule
|string
a|Schedule associated with volume.

|space_savings
|link:#space_savings[space_savings]
a|

|state
|string
a|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
a|Storage efficiency mode used by volume. This parameter is supported only on AFF platform.

|type
|string
a|Sis Type of the volume.

|volume_path
|string
a|Absolute volume path of the volume.

|===

[#status]
[.api-collapsible-fifth-title]
status

[cols=3*,options=header]
|===
|Name
|Type
|Description

```
|code
|string
a|Encryption progress message code.
```

```
|message
|string
a|Encryption progress message.
```

```
|===
```

```
[#encryption]
[.api-collapsible-fifth-title]
encryption
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|enabled
|boolean
```

```
a|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
a|Encryption key creation time of the volume.
```

```
|key_id
|string
a|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
a|Specifies an additional key manager attribute that is an identifier-
value pair, separated by '='. For example, CRN=unique-value. This
parameter is required when using the POST method and an IBM Key Lore key
manager is configured on the SVM.

|rekey
|boolean
a|If set to 'true', re-encrypts the volume with a new key. Valid in PATCH.

|state
|string
a|Volume encryption state.
encrypted &dash; The volume is completely encrypted.
encrypting &dash; Encryption operation is in progress.
partial &dash; Some constituents are encrypted and some are not.
Applicable only for FlexGroup volume.
rekeying. Encryption of volume with a new key is in progress.
unencrypted &dash; The volume is a plain-text one.

|status
|link:#status[status]
a|

|type
|string
a|Volume encryption type.
none &dash; The volume is a plain-text one.
volume &dash; The volume is encrypted with NVE (NetApp Volume Encryption).
aggregate &dash; The volume is encrypted with NAE (NetApp Aggregate
Encryption).

|===

[#error_state]
[.api-collapsible-fifth-title]
error_state

[cols=3*,options=header]
|===
|Name
|Type

```

|Description

|has_bad_blocks

|boolean

a|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

a|Indicates whether the file system has any inconsistencies.
true ‐ File system is inconsistent.
false ‐ File system is not inconsistent.

|===

[#files]

[.api-collapsible-fifth-title]

files

[cols=3*,options=header]

|===

|Name

|Type

|Description

|maximum

|integer

a|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

a|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]  
[.api-collapsible-fifth-title]  
flash_pool
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|cache_eligibility
```

```
|string
```

a|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
```

a|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
```

a|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]  
[.api-collapsible-fifth-title]  
flexgroup
```

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|name
|string
a|Name of the FlexGroup volume that the constituent is part of.

|uuid
|string
a|Unique identifier for the FlexGroup volume that the constituent is part
of.

|===

[#guarantee]
[.api-collapsible-fifth-title]
guarantee

[cols=3*,options=header]
|===
|Name
|Type
|Description

|honored
|boolean
a|Is the space guarantee of this volume honored in the aggregate?

|type
|string
a|The type of space guarantee of this volume in the aggregate.

|===

[#iops]
[.api-collapsible-fifth-title]
iops

The rate of I/O operations observed at the storage object.

```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
```

```
[.api-collapsible-fifth-title]
```

```
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be
```


metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#cloud]
[.api-collapsible-fifth-title]
cloud
```

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|duration
|string
a|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
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#flexcache]
[.api-collapsible-fifth-title]
flexcache
```

Performance number for FlexCache used to measure cache effectiveness.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|bandwidth_savings
```

```
|integer
```

```
a|Bandwidth savings denoting the amount of data served locally by the cache, in bytes.
```

```
|cache_miss_percent
|integer
a|Cache miss percentage.
```

```
|duration
|string
a|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
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

Performance numbers, such as IOPS, latency and throughput.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|cloud
|link:#cloud[cloud]
a|Performance numbers (IOPS and latency) for cloud store. These numbers
are relevant only for volumes hosted on FabricPools.
```

|duration
|string
a|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
|link:#flexcache[flexcache]
a|Performance number for FlexCache used to measure cache effectiveness.

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|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
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

```
|===
```

```
[#movement]
```

```
[.api-collapsible-fifth-title]
```

```
movement
```

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|cutover_window
```

```
|integer
```

a|Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.

```
|destination_aggregate
```

```
|link:#destination_aggregate[destination_aggregate]
```

a|Aggregate

```
|percent_complete
```

```
|integer
```

a|Completion percentage

```
|start_time
```

```
|string
```

a|Start time of volume move.

```
|state
```

```
|string
```

a|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
a|Tiering policy for FabricPool
```

```
|===
```

```
[#export_policy]
[.api-collapsible-fifth-title]
export_policy
```

Export Policy

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|id
|integer
a|
```

```
|name
|string
a|
```

```
|===
```

```
[#junction_parent]
[.api-collapsible-fifth-title]
junction_parent
```

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|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
a|Unique identifier for the parent volume.

|===

[#nas]
[.api-collapsible-fifth-title]
nas

[cols=3*,options=header]
|===
|Name
|Type
|Description

|export_policy
|link:#export_policy[export_policy]
a|Export Policy

|gid
|integer
a|The UNIX group ID of the volume. Valid in POST or PATCH.

|junction_parent
|link:#junction_parent[junction_parent]
a|

```


|path

|string

a|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

a|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

a|The UNIX user ID of the volume. Valid in POST or PATCH.

|unix_permissions

|integer

a|UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth digit), first digit is assumed to be zero.

|===

[#policy]

[.api-collapsible-fifth-title]

policy

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|max_throughput_iops

|integer

a|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

a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

|min_throughput_iops

|integer

a|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

a|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

a|The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.

|uuid

|string

a|The QoS policy group UUID. This is mutually exclusive with name and

other QoS attributes during POST and PATCH.

|===

[#qos]

[.api-collapsible-fifth-title]

qos

QoS information

[cols=3*,options=header]

|===

|Name

|Type

|Description

|policy

|link:#policy[policy]

a|

|===

[#quota]

[.api-collapsible-fifth-title]

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|enabled

|boolean

a|This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVols or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVols or FlexGroup volumes when the quota state is "off".

```
|state
|string
a|Quota state of the volume
```

```
|===
```

```
[#last_error]
[.api-collapsible-fifth-title]
last_error
```

Error information for the last failed file move on the constituent.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
```

```
|integer
```

a|Error code of the last file move error on the constiuent.

```
|destination
```

```
|integer
```

a|DSID of the destination constituent of the last file move error on the constiuent.

```
|file_id
```

```
|integer
```

a|File ID of the last file move error on the constiuent.

```
|time
```

```
|string
```

a|Time of the last file move error on the constiuent.

```
|===
```

```
[#movement]
```

```
[.api-collapsible-fifth-title]
```

```
movement
```

Properties on this constituent related to file movement.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|file_moves_started
```

```
|integer
```

```
a|Number of file moves started on this constituent.
```

```
|last_error
```

```
|link:#last_error[last_error]
```

```
a|Error information for the last failed file move on the constituent.
```

```
|most_recent_start_time
```

```
|string
```

```
a|Start time of the most recent file move on the constiuent.
```

```
|===
```

```
[#blocks_skipped]
```

```
[.api-collapsible-fifth-title]
```

```
blocks_skipped
```

Number of blocks skipped by the scanner on this constiuent due to various reasons.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|efficiency_blocks
```

```
|integer
```

```
a|Number of blocks skipped by the scanner on this constituent because
```

storage efficiency lost, in blocks, would be too high.

|efficiency_percent

|integer

a|Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.

|fast_truncate

|integer

a|Number of blocks skipped by the scanner on this constituent because fast truncate is currently running on files.

|footprint_invalid

|integer

a|Number of blocks skipped by the scanner on this constituent because of files with invalid space footprints.

|in_snapshot

|integer

a|Number of blocks skipped by the scanner on this constituent because of files in Snapshot copies.

|incompatible

|integer

a|Number of blocks skipped by the scanner on this constituent because of incompatible files.

|metadata

|integer

a|Number of blocks skipped by the scanner on this constituent because of metadata files.

|on_demand_destination

|integer

a|Number of blocks skipped by the scanner on this constituent because of on demand destination files.

|other

|integer

a|Number of blocks skipped by the scanner on this constituent for all other reasons.

|remote_cache

|integer

a|Number of blocks skipped by the scanner on this constituent because of remote caches.

|too_large

|integer

a|Number of blocks skipped by the scanner on this constituent because of files that are larger than `rebalancing.max_file_size`.

|too_small

|integer

a|Number of blocks skipped by the scanner on this constituent because of files that are smaller than `rebalancing.min_file_size`.

|write_fenced

|integer

a|Number of blocks skipped by the scanner on this constituent because of files fenced for write operations.

|===

[#files_skipped]

[.api-collapsible-fifth-title]

files_skipped

Number of files skipped by the scanner on this constituent due to various reasons.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|efficiency_blocks

|integer

a|Number of files skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.

|efficiency_percent

|integer

a|Number of files skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.

|fast_truncate

|integer

a|Number of files skipped by the scanner on this constituent because fast truncate is currently running on the file.

|footprint_invalid

|integer

a|Number of files skipped by the scanner on this constituent because their space footprints are invalid.

|in_snapshot

|integer

a|Number of files skipped by the scanner on this constituent because they are trapped in Snapshot copies.

|incompatible

|integer

a|Number of files skipped by the scanner on this constituent because they are incompatible.

|metadata

|integer

a|Number of files skipped by the scanner on this constituent because they metadata files.

|on_demand_destination

|integer

a|Number of files skipped by the scanner on this constituent because they are on demand destinations.

|other


```

|integer
a|Number of files skipped by the scanner on this constituent for all other
reasons.

|remote_cache
|integer
a|Number of files skipped by the scanner on this constituent because they
are remote caches.

|too_large
|integer
a|Number of files skipped by the scanner on this constituent because they
are larger than rebalancing.max_file_size.

|too_small
|integer
a|Number of files skipped by the scanner on this constituent because they
are smaller than rebalancing.min_file_size.

|write_fenced
|integer
a|Number of files skipped by the scanner on this constituent because they
are fenced for write operations.

|===

[#scanner]
[.api-collapsible-fifth-title]
scanner

Properties related to determining which files to move and where to move
them to.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|blocks_scanned

```

```
|integer
a|Number of blocks scanned on this constituent.
```

```
|blocks_skipped
|link:#blocks_skipped[blocks_skipped]
a|Number of blocks skipped by the scanner on this constiuent due to
various reasons.
```

```
|files_scanned
|integer
a|Number of files scanned on this constituent.
```

```
|files_skipped
|link:#files_skipped[files_skipped]
a|Number of files skipped by the scanner on this constiuent due to
various reasons.
```

```
|===
```

```
[#engine]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|movement
|link:#movement[movement]
a|Properties on this constituent related to file movement.
```

```
|scanner
```

```

|link:#scanner[scanner]
a|Properties related to determining which files to move and where to move
them to.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code

```

```
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
[#rebalancing]
[.api-collapsible-fifth-title]
rebalancing
```

Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|data_moved
|integer
a|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
|link:#engine[engine]
a|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

a|Specifies whether or not to exclude files that are stuck in Snapshot copies 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

a|Represents the percentage the volume is out of balance.

|imbalance_size

|integer

a|Represents how much the volume is out of balance, in bytes.

|max_constituent_imbalance_percent

|integer

a|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

a|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

a|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>+++

|max_threshold

|integer

a|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

a|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.

|min_threshold

|integer

a|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[link:#error[error]]

a|Capacity rebalancing notice messages.

|runtime

|string

a|Duration the capacity rebalancing operation has been running.

|start_time

|string

a|Time when the current capacity rebalancing operation started, or when a future scheduled rebalancing operation begins.

|state

|string

a|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.

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.

idle ‐ capacity rebalancing is running on the constituent, however, no active scanning or file movement is currently occurring.

scanning ‐ the constituent's file system is being scanned to find files to move and determine free space.

rebalancing_source ‐ a file is being moved off of the constituent.

rebalancing_dest ‐ a file is being moved to the constituent.

not_running ‐ capacity rebalancing is not running on the constituent.

unknown ‐ the system was unable to determine the rebalancing state for the constituent.

|stop_time

|string

a|Time when the capacity rebalancing operation stopped.

|target_used

|integer

a|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

a|Represents the used size of each constituent, as determined by the rebalancing engine. Calculated by subtracting the size used by Snapshot copies, the size of files pending deletion and the size of filesystem metadata from the volume used size.

|===

[#retention]

[.api-collapsible-fifth-title]

retention

[cols=3*,options=header]

|===

|Name

|Type

|Description

|default

|string

a|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>+++

|maximum

|string

a|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>+++

|minimum

|string

a|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]
[.api-collapsible-fifth-title]
snaplock
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|append_mode_enabled
```

```
|boolean
```

a|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.

```
|autocommit_period
```

```
|string
```

a|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

a|This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.

|expiry_time

|string

a|Expiry time of the volume.

|is_audit_log

|boolean

a|Indicates if this volume has been configured as SnapLock audit log volume for the SVM .

|litigation_count

|integer

a|Litigation count indicates the number of active legal-holds on the volume.

|privileged_delete

|string

a|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

|link:#retention[retention]

a|

|type

```
|string
a|The SnapLock type of the volume.
compliance &dash; A SnapLock Compliance (SLC) volume provides the highest
level of WORM protection and an administrator cannot destroy a SLC volume
if it contains unexpired WORM files.
enterprise &dash; An administrator can delete a SnapLock Enterprise (SLE)
volume.
non_snaplock &dash; Indicates the volume is non-snaplock.
```

```
|unspecified_retention_file_count
|integer
a|Indicates the number of files with an unspecified retention time in the
volume.
```

```
|===
```

```
[#destinations]
[.api-collapsible-fifth-title]
destinations
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|is_cloud
|boolean
a|Specifies whether a volume is a SnapMirror source volume, using
SnapMirror to protect its data to a cloud destination.
```

```
|is_ontap
|boolean
a|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]
[.api-collapsible-fifth-title]
snapmirror
```

Specifies attributes for SnapMirror protection.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|destinations
|link:#destinations[destinations]
a|
```

```
|is_protected
|boolean
a|Specifies whether a volume is a SnapMirror source volume, using
SnapMirror to protect its data.
```

```
|===
```

```
[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy
```

This is a reference to the Snapshot copy policy.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#logical_space]
[.api-collapsible-fifth-title]
logical_space
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|available
```

```
|integer
```

```
a|The amount of space available in this volume with storage efficiency
space considered used, in bytes.
```

```
|enforcement
```

```
|boolean
```

```
a|Specifies whether space accounting for operations on the volume is done
along with storage efficiency.
```

```
|reporting
```

```
|boolean
```

```
a|Specifies whether space reporting on the volume is done along with
storage efficiency.
```

```
|used
```

```
|integer
```

```
a|SUM of (physical-used, shared_refs, compression_saved_in_plane0,
vbn_zero, future_blk_cnt), in bytes.
```

```
|used_by_afs
```

```
|integer
```

```
a|The virtual space used by AFS alone (includes volume reserves) and along
with storage efficiency, in bytes.
```

|used_by_snapshots
|integer
a|Size that is logically used across all Snapshot copies in the volume, in bytes.

|used_percent
|integer
a|SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), as a percentage.

|===

[#autodelete]
[.api-collapsible-fifth-title]
autodelete

[cols=3*,options=header]

|===

|Name
|Type
|Description

|commitment

|string

a|By default, Snapshot copy autodelete does not delete Snapshot copies locked by Snapmirror, clones of a volume, a LUN, an NVMe namespace, or a file. Deletion of Snapshot copies locked by these applications is specified using this option. The default value is try.

|defer_delete

|string

a|Allows the user to inform Snapshot copy autodelete to defer the deletion of a specified Snapshot copy until the end. The default value is user_created.

|delete_order

|string

a|Specifies the order in which Snapshot copy autodelete occurs. Ordering is done using the date and time the Snapshot copy is created. The default value is oldest_first.

|enabled
|boolean
a|Specifies whether Snapshot copy autodelete is currently enabled on this volume.

|prefix
|string
a|Specifies the prefix of the Snapshot copy which if matched, is deleted last. Used with autodelete_defer_delete when used with a prefix value.

|target_free_space
|integer
a|Snapshot copies 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
a|Specifies when the system should trigger an autodelete of Snapshot copies. When set to `_volume_`, autodelete is triggered based on volume fullness. When set to `_snap_reserve_`, autodelete is triggered based on Snapshot copy reserve fullness. The default value is `_volume_`.

|===

[#snapshot]
[.api-collapsible-fifth-title]
snapshot

[cols=3*,options=header]

|===

|Name
|Type
|Description

|autodelete
|link:#autodelete[autodelete]
a|

|autodelete_enabled
|boolean

a|Specifies whether Snapshot copy 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

a|Specifies when the system should trigger an autodelete of Snapshot copies. When set to `_volume_`, autodelete is triggered based on volume fullness. When set to `_snap_reserve_`, autodelete is triggered based on Snapshot copy 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

a|Size available for Snapshot copies within the Snapshot copy reserve, in bytes.

|reserve_percent

|integer

a|The space that has been set aside as a reserve for Snapshot copy usage, in percent.

|reserve_size

|integer

a|Size in the volume that has been set aside as a reserve for Snapshot copy usage, in bytes.

|space_used_percent

|integer

a|Percentage of snapshot reserve size that has been used.

|used

|integer

a|The total space used by Snapshot copies in the volume, in bytes.

|===

[#space]

[.api-collapsible-fifth-title]

space

[cols=3*,options=header]

|===

|Name

|Type

|Description

|afs_total

|integer

a|Total size of AFS, excluding snap-reserve, in bytes.

|auto_adaptive_compression_footprint_data_reduction

|integer

a|Savings achieved due to Auto Adaptive Compression, in bytes.

|available

|integer

a|The available space, in bytes.

|available_percent

|integer

a|The space available, as a percent.

|block_storage_inactive_user_data

|integer

a|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

a|Percentage of size that is physically used in the performance tier of the volume.

|capacity_tier_footprint

|integer

a|Space used by capacity tier for this volume in the FabricPool aggregate, in bytes.

|capacity_tier_footprint_data_reduction
|integer
a|Savings achieved in the space used by the capacity tier for this volume in the FabricPool aggregate, in bytes.

|cross_volume_dedupe_metafiles_footprint
|integer
a|Cross volume deduplication metadata footprint, in bytes.

|cross_volume_dedupe_metafiles_temporary_footprint
|integer
a|Cross volume temporary deduplication metadata footprint, in bytes.

|dedupe_metafiles_footprint
|integer
a|Deduplication metadata footprint, in bytes.

|dedupe_metafiles_temporary_footprint
|integer
a|Temporary deduplication metadata footprint, in bytes.

|delayed_free_footprint
|integer
a|Delayed free blocks footprint, in bytes.

|effective_total_footprint
|integer
a|Volume footprint after efficiency savings, in bytes.

|expected_available
|integer
a|Size that should be available for the volume, irrespective of available size in the aggregate, in bytes.

|file_operation_metadata
|integer
a|File operation metadata footprint, in bytes.

|filesystem_size
|integer
a|Total usable size of the volume, in bytes.

|filesystem_size_fixed
|boolean
a|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
a|Data used for this volume in the aggregate, in bytes.

|fractional_reserve
|integer
a|Used to change the amount of space reserved for overwrites of reserved objects in a volume.

|full_threshold_percent
|integer
a|Volume full threshold percentage at which EMS warnings can be sent.

|is_used_stale
|boolean
a|Specifies if the virtual space used is stale.

|large_size_enabled
|boolean
a|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
a|Space used by the local tier for this volume in the aggregate, in bytes.

```
|logical_space
|link:#logical_space[logical_space]
a|

|metadata
|integer
a|Space used by the volume metadata in the aggregate, in bytes.

|nearly_full_threshold_percent
|integer
a|Volume nearly full threshold percentage at which EMS warnings can be sent.

|over_provisioned
|integer
a|The amount of space not available for this volume in the aggregate, in bytes.

|overwrite_reserve
|integer
a|Reserved space for overwrites, in bytes.

|overwrite_reserve_used
|integer
a|Overwrite logical reserve space used, in bytes.

|percent_used
|integer
a|Percentage of the volume size that is used.

|performance_tier_footprint
|integer
a|Space used by the performance tier for this volume in the FabricPool aggregate, in bytes.

|physical_used
|integer
a|Size that is physically used in the volume, in bytes.
```

```
|physical_used_percent
|integer
a|Size that is physically used in the volume, as a percentage.

|size
|integer
a|Total provisioned size. The default size is equal to the minimum size of
20MB, in bytes.

|size_available_for_snapshots
|integer
a|Available space for Snapshot copies from snap-reserve, in bytes.

|snapmirror_destination_footprint
|integer
a|SnapMirror destination footprint, in bytes.

|snapshot
|link:#snapshot[snapshot]
a|

|snapshot_reserve_unusable
|integer
a|Snapshot reserve that is not available for Snapshot copy creation, in
bytes.

|snapshot_spill
|integer
a|Space used by the snapshot copies beyond the snap-reserve, in bytes.

|total_footprint
|integer
a|Data and metadata used for this volume in the aggregate, in bytes.

|used
|integer
a|The virtual space used (includes volume reserves) before storage
efficiency, in bytes.
```

```
|used_by_afs
|integer
a|The space used by Active Filesystem, in bytes.
```

```
|user_data
|integer
a|User data, in bytes.
```

```
|volume_guarantee_footprint
|integer
a|Space reserved for future writes in the volume, in bytes.
```

```
|===
```

```
[#access]
[.api-collapsible-fifth-title]
access
```

Raw count and latency data for access operations.

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|count
|integer
a|Number of operations of the given type performed on this volume.
```

```
|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
audit
```

Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|count
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
|integer
```

a|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]
[.api-collapsible-fifth-title]
dir
```

Raw count and latency data for directory-create operations.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```



```
|count
|integer
a|Number of operations of the given type performed on this volume.

|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
file
```

Raw count and latency data for file-create operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|count
|integer
a|Number of operations of the given type performed on this volume.
```

```
|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
other
```

Raw count and latency data for create operations on objects other than files, directories and symlinks.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|count
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
|integer
```

a|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]
[.api-collapsible-fifth-title]
symlink
```

Raw count and latency data for symlink-create operations.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|count
|integer
a|Number of operations of the given type performed on this volume.

|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
create
```

Raw count and latency data for create operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|dir
```

```
|link:#dir[dir]
```

a|Raw count and latency data for directory-create operations.

```
|file
```

```
|link:#file[file]
```

a|Raw count and latency data for file-create operations.

```
|other
```

```
|link:#other[other]
```

a|Raw count and latency data for create operations on objects other than files, directories and symlinks.

```
|symlink
|link:#symlink[symlink]
a|Raw count and latency data for symlink-create operations.
```

```
|===
```

```
[#getattr]
[.api-collapsible-fifth-title]
getattr
```

Raw count and latency data for getattr operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
```

a|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]
[.api-collapsible-fifth-title]
link
```

Raw count and latency data for link operations.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|count
|integer
a|Number of operations of the given type performed on this volume.

|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
lock

Raw count and latency data for lock operations.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|Number of operations of the given type performed on this volume.

|total_time
|integer
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
lookup
```

Raw count and latency data for lookup operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
```

a|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]
```

```
[.api-collapsible-fifth-title]
```

```
open
```

Raw count and latency data for open operations.

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|count
|integer
a|Number of operations of the given type performed on this volume.

|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
read

Raw count and latency data for read operations, including histograms
categorizing operations by size and latency.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|Number of operations of the given type performed on this volume.

|total_time
|integer
a|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]
a|
```

```
|volume_protocol_latency_histogram_labels
|array[string]
a|Labels for the latency histogram, ranging from <2us to >20s.
```

```
|volume_protocol_size_histogram_counts
|array[integer]
a|
```

```
|volume_protocol_size_histogram_labels
|array[string]
a|Labels for the size histogram, ranging from <4KB to >1024KB.
```

```
|===
```

```
[#readdir]
[.api-collapsible-fifth-title]
readdir
```

Raw count and latency data for readdir operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|count
|integer
a|Number of operations of the given type performed on this volume.
```

```
|total_time
|integer
```


a|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]
[.api-collapsible-fifth-title]
readlink

Raw count and latency data for readlink operations.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|count
|integer

a|Number of operations of the given type performed on this volume.

|total_time
|integer

a|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]
[.api-collapsible-fifth-title]
rename

Raw count and latency data for rename operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
setattr
```

Raw count and latency data for setattr operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
[.api-collapsible-fifth-title]
unlink
```

Raw count and latency data for unlink operations.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
```

a|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]
[.api-collapsible-fifth-title]
watch
```

Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
write
```

Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

a|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]

a|

|volume_protocol_latency_histogram_labels

|array[string]

a|Labels for the latency histogram, ranging from <2us to >20s.

|volume_protocol_size_histogram_counts

|array[integer]

a|

|volume_protocol_size_histogram_labels

|array[string]

a|Labels for the size histogram, ranging from <4KB to >1024KB.

|===

[#cifs_ops_raw]

[.api-collapsible-fifth-title]

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.

[cols=3*,options=header]

|===

|Name

|Type

|Description

```
|access
|link:#access[access]
a|Raw count and latency data for access operations.

|audit
|link:#audit[audit]
a|Raw count and latency data for audit operations. These statistics are
only applicable for CIFS protocol operations.

|create
|link:#create[create]
a|Raw count and latency data for create operations.

|getattr
|link:#getattr[getattr]
a|Raw count and latency data for getattr operations.

|link
|link:#link[link]
a|Raw count and latency data for link operations.

|lock
|link:#lock[lock]
a|Raw count and latency data for lock operations.

|lookup
|link:#lookup[lookup]
a|Raw count and latency data for lookup operations.

|open
|link:#open[open]
a|Raw count and latency data for open operations.

|read
|link:#read[read]
a|Raw count and latency data for read operations, including histograms
categorizing operations by size and latency.
```

```
|readdir
|link:#readdir[readdir]
a|Raw count and latency data for readdir operations.
```

```
|readlink
|link:#readlink[readlink]
a|Raw count and latency data for readlink operations.
```

```
|rename
|link:#rename[rename]
a|Raw count and latency data for rename operations.
```

```
|setattr
|link:#setattr[setattr]
a|Raw count and latency data for setattr operations.
```

```
|unlink
|link:#unlink[unlink]
a|Raw count and latency data for unlink operations.
```

```
|watch
|link:#watch[watch]
a|Raw count and latency data for watch operations. These statistics are
only applicable for CIFS protocol operations.
```

```
|write
|link:#write[write]
a|Raw count and latency data for write operations, including histograms
categorizing operations by size and latency.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
```

```
[.api-collapsible-fifth-title]
```

```
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```



```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#cloud]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|iops_raw
|link:#iops_raw[iops_raw]
a|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
|link:#latency_raw[latency_raw]
a|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
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#flexcache_raw]
[.api-collapsible-fifth-title]
flexcache_raw
```

Performance numbers for FlexCache used to measure cache effectiveness.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|cache_miss_blocks
|integer
```

a|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
a|Total blocks requested by the client.
```

```
|status
|string
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#nfs_ops_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|access
|link:#access[access]
a|Raw count and latency data for access operations.

|audit
|link:#audit[audit]
a|Raw count and latency data for audit operations. These statistics are
only applicable for CIFS protocol operations.

|create
|link:#create[create]
a|Raw count and latency data for create operations.

|getattr
|link:#getattr[getattr]
a|Raw count and latency data for getattr operations.

|link
|link:#link[link]
a|Raw count and latency data for link operations.

|lock
|link:#lock[lock]
a|Raw count and latency data for lock operations.

|lookup
|link:#lookup[lookup]
a|Raw count and latency data for lookup operations.

|open
|link:#open[open]
a|Raw count and latency data for open operations.

|read
|link:#read[read]
a|Raw count and latency data for read operations, including histograms
categorizing operations by size and latency.
```

```
|readdir
|link:#readdir[readdir]
a|Raw count and latency data for readdir operations.
```

```
|readlink
|link:#readlink[readlink]
a|Raw count and latency data for readlink operations.
```

```
|rename
|link:#rename[rename]
a|Raw count and latency data for rename operations.
```

```
|setattr
|link:#setattr[setattr]
a|Raw count and latency data for setattr operations.
```

```
|unlink
|link:#unlink[unlink]
a|Raw count and latency data for unlink operations.
```

```
|watch
|link:#watch[watch]
a|Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.
```

```
|write
|link:#write[write]
a|Raw count and latency data for write operations, including histograms categorizing operations by size and latency.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

```
[.api-collapsible-fifth-title]
```

```
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

|cifs_ops_raw

|link:#cifs_ops_raw[cifs_ops_raw]

a|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

|link:#cloud[cloud]

a|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

|link:#flexcache_raw[flexcache_raw]

a|Performance numbers for FlexCache used to measure cache effectiveness.

|iops_raw

|link:#iops_raw[iops_raw]

a|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

|link:#latency_raw[latency_raw]

a|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

|link:#nfs_ops_raw[nfs_ops_raw]

a|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

a|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
|link:#throughput_raw[throughput_raw]
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

SVM containing the volume. Required on POST.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
 |_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```



```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#tiering]
[.api-collapsible-fifth-title]
tiering
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|min_cooling_days
```

```
|integer
```

a|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]
```

a|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".

```
|policy
```

```
|string
```

a|Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a

single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH. all ‐ This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.

auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store

none ‐ Volume blocks will not be tiered to the cloud store.

snapshot_only ‐ This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days.

|storage_class

|string

a|This parameter specifies the storage class that a FabricPool uses. This feature is only available on volumes in a FabricPools on FSx or Cloud Volumes ONTAP for AWS.

* Default value: 1

* enum: ["default", "S3_standard", "S3_standard_IA", "S3_glacier_IR"]

* Introduced in: 9.13

* x-ntap-modifyOnly: true

* x-nullable: true

|supported

|boolean

a|This parameter specifies whether or not FabricPools are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only FabricPool aggregates are used if this parameter is set to true and only non FabricPool aggregates are used if this parameter is set to false. Tiering support for a FlexGroup can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only does the volume need to support tiering by using FabricPools, the tiering "policy" must not be 'none'. A volume that uses FabricPools but has a tiering "policy" of 'none' supports tiering, but will not tier any data.

|===

[#volume]

[.api-collapsible-fifth-title]

volume

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|_tags

|array[string]

a|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

a|Indicates whether or not access time updates are enabled on the volume.

|activity_tracking

|link:#activity_tracking[activity_tracking]

a|

|aggregates

|array[link:#aggregates[aggregates]]

a|Aggregate hosting the volume. Required on POST.

|aggressive_readahead_mode

|string

a|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.

* Default value: 1

* enum: ["none", "file_prefetch"]

* Introduced in: 9.13

* x-nullable: true

|analytics

```
|link:#analytics[analytics]
```

```
a|
```

```
|anti_ransomware
```

```
|link:#anti_ransomware[anti_ransomware]
```

```
a|Anti-ransomware related information of the volume.
```

```
|anti_ransomware_state
```

```
|string
```

```
a|The Anti-ransomware state of the volume. If no "anti_ransomware_state" property is specified, the volume inherits the value from its parent SVM's "anti_ransomware_default_volume_state" property. If this value is "disabled", Anti-ransomware is disabled on the volume. If this value is "enabled", Anti-ransomware is enabled on the volume and alerts are raised if any suspect is detected for those volumes. If this value is "dry_run", Anti-ransomware is enabled in the dry-run or learning mode on the volume. The "dry_run" state is same as the "enabled" state except that the analytics data is used here for learning. No alerts are raised for any detections or violations. If this value is "paused", Anti-ransomware is paused on the volume. Additionally, three more states are available, which are only valid for GET. If this value is "disable_in_progress", Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. If this value is "enable_paused", Anti-ransomware is paused on the volume from its earlier enabled state. If this value is "dry_run_paused", Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. For POST, the valid Anti-ransomware states are only "disabled", "enabled" and "dry_run", whereas for PATCH, "paused" is also valid along with the three valid states for POST.
```

```
|application
```

```
|link:#application[application]
```

```
a|
```

```
|asynchronous_directory_delete
```

```
|link:#asynchronous_directory_delete[asynchronous_directory_delete]
```

```
a|Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.
```

```
|autosize
```

```
|link:#autosize[autosize]
```

```
a|
```

```
|clone
```

|link:#clone[clone]

a|

|cloud_retrieval_policy

|string

a|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

a|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 on FSx or Cloud Volumes ONTAP.

* Introduced in: 9.13

* x-ntap-readModify: true

* x-nullable: true

|comment

|string

a|A comment for the volume. Valid in POST or PATCH.

|consistency_group

|link:#consistency_group[consistency_group]

a|Consistency group the volume is part of.

a|

|flash_pool
|link:#flash_pool[flash_pool]
a|

|flexcache_endpoint_type
|string
a|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
|link:#flexgroup[flexgroup]
a|

|granular_data
|boolean
a|State of granular data on the volume. This setting is true by default when creating a new FlexGroup volume, but can be specified as false at the time of creation via a POST request. On FlexVol volumes, the setting is always false, as only FlexGroup volumes and FlexGroup constituents support this feature. Once enabled, this setting can only be disabled by restoring a Snapshot copy. Earlier versions of ONTAP (pre 9.11) 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 copy that was taken before the setting was enabled.

* Introduced in: 9.12
* x-nullable: true

|guarantee
|link:#guarantee[guarantee]
a|

|is_object_store
|boolean
a|Specifies whether the volume is provisioned for an object store server.

|is_svm_root
|boolean
a|Specifies whether the volume is a root volume of the SVM it belongs to.

|language
|string
a|Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.

|max_dir_size
|integer
a|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
|link:#metric[metric]
a|Performance numbers, such as IOPS, latency and throughput.

|movement
|link:#movement[movement]
a|Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

|msid
|integer
a|The volume's Master Set ID.

|name
|string
a|Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.

|nas
|link:#nas[nas]
a|

|optimize_aggregates

|boolean
a|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
|link:#qos[qos]
a|QoS information

|queue_for_encryption
|boolean
a|Specifies whether the volume is queued for encryption.

|quota
|link:#quota[quota]
a|Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

|rebalancing
|link:#rebalancing[rebalancing]
a|Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

|scheduled_snapshot_naming_scheme
|string
a|Naming Scheme for automatic Snapshot copies:

* create_time - Automatic Snapshot copies are saved as per the start of their current date and time.

* ordinal - Latest automatic snapshot copy is saved as +++<scheduled_frequency>+++0 and subsequent copies will follow the create_time naming convention.+++</scheduled_frequency>+++

|size
|integer
a|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
|link:#snaplock[snaplock]
a|

|snapmirror
|link:#snapmirror[snapmirror]
a|Specifies attributes for SnapMirror protection.

|snapshot_count
|integer
a|Number of Snapshot copies in the volume.

|snapshot_directory_access_enabled
|boolean
a|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
a|Specifies whether or not snapshot copy locking is enabled on the volume.

|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the Snapshot copy policy.

|space
|link:#space[space]
a|

|state
|string
a|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

|link:#statistics[statistics]

a|These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

|status

|array[string]

a|Describes the current status of a volume.

|style

|string

a|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 constituents.

|svm

|link:#svm[svm]

a|SVM containing the volume. Required on POST.

|tiering

|link:#tiering[tiering]

a|

|type

|string

a|Type of the volume.

rw ‐ read-write volume.

dp ‐ data-protection volume.
ls ‐ load-sharing `dp` volume. Valid in GET.

|use_mirrored_aggregates
|boolean

a|Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.

|uuid
|string

a|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

|===

//end collapsible .Definitions block
=====

[[ID27ccb59b3432034445f7f94ee4b6858d]]
= Create a volume on an SVM and storage aggregates

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/storage/volumes`#

Introduced In: 9.6

Creates a volume on a specified SVM and storage aggregates.

== Required properties

```
* `svm.uuid` or `svm.name` - Existing SVM in which to create the volume.
* `name` - Name of the volume.
* `aggregates.name` or `aggregates.uuid` - Existing aggregates in which to
create the volume.
```

== Default property values

```
* `state` - _online_
* `size` - _20MB_
* `style` - _flexvol_
* `type` - _rw_
* `encryption.enabled` - _false_
* `snapshot_policy.name` - _default_
* `guarantee.type` - _volume_
* `anti_ransomware.state` - _default_
```

== Related ONTAP commands

```
* `volume create`
* `volume clone create`
```

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|return_timeout
|integer
|query
|False
```

a|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

a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|_tags

|array[string]

a|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

a|Indicates whether or not access time updates are enabled on the volume.

|activity_tracking

|link:#activity_tracking[activity_tracking]

a|

|aggregates

|array[link:#aggregates[aggregates]]

a|Aggregate hosting the volume. Required on POST.

|aggressive_readahead_mode
|string
a|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.

* Default value: 1
* enum: ["none", "file_prefetch"]
* Introduced in: 9.13
* x-nullable: true

|analytics
|link:#analytics[analytics]
a|

|anti_ransomware
|link:#anti_ransomware[anti_ransomware]
a|Anti-ransomware related information of the volume.

|anti_ransomware_state
|string
a|The Anti-ransomware state of the volume. If no "anti_ransomware_state" property is specified, the volume inherits the value from its parent SVM's "anti_ransomware_default_volume_state" property. If this value is "disabled", Anti-ransomware is disabled on the volume. If this value is "enabled", Anti-ransomware is enabled on the volume and alerts are raised if any suspect is detected for those volumes. If this value is "dry_run", Anti-ransomware is enabled in the dry-run or learning mode on the volume. The "dry_run" state is same as the "enabled" state except that the analytics data is used here for learning. No alerts are raised for any detections or violations. If this value is "paused", Anti-ransomware is paused on the volume. Additionally, three more states are available, which are only valid for GET. If this value is "disable_in_progress", Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. If this value is "enable_paused", Anti-ransomware is paused on the volume from its earlier enabled state. If this value is "dry_run_paused", Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. For POST, the valid Anti-ransomware states are only "disabled", "enabled" and "dry_run", whereas for PATCH, "paused" is also valid along with the three valid states for POST.

|application

|link:#application[application]

a|

|asynchronous_directory_delete

|link:#asynchronous_directory_delete[asynchronous_directory_delete]

a|Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.

|autosize

|link:#autosize[autosize]

a|

|clone

|link:#clone[clone]

a|

|cloud_retrieval_policy

|string

a|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

a|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 on FSx or Cloud Volumes ONTAP.


```
* Introduced in: 9.13
* x-ntap-readModify: true
* x-nullable: true
```

```
|comment
|string
a|A comment for the volume. Valid in POST or PATCH.
```

```
|consistency_group
|link:#consistency_group[consistency_group]
a|Consistency group the volume is part of.
```

```
|constituents
|array[link:#constituents[constituents]]
a|FlexGroup Constituents. FlexGroup 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
a|Specifies the number of times to iterate over the aggregates listed with
the "aggregates.name" or "aggregates.uuid" when creating or expanding a
FlexGroup volume. If a volume is being created on a single aggregate, the
system creates a flexible volume if the "constituents_per_aggregate" field
is not specified, or a FlexGroup volume if it is specified. If a volume is
being created on multiple aggregates, the system always creates a
FlexGroup volume. The root constituent of a FlexGroup volume is always
placed on the first aggregate in the list, unless 'optimize_aggregates' is
specified as 'true'.
```

```
|convert_unicode
|boolean
a|Specifies whether directory Unicode format conversion is enabled when
directories are accessed by NFS clients.
```

```
|create_time
|string
a|Creation time of the volume. This field is generated when the volume is
created.
```

```
|efficiency
|link:#efficiency[efficiency]
a|
```

```
|encryption
|link:#encryption[encryption]
a|
```

```
|error_state
|link:#error_state[error_state]
a|
```

```
|files
|link:#files[files]
a|
```

```
|flash_pool
|link:#flash_pool[flash_pool]
a|
```

```
|flexcache_endpoint_type
|string
a|FlexCache endpoint type.
none &dash; The volume is neither a FlexCache nor origin of any FlexCache.
cache &dash; The volume is a FlexCache volume.
origin &dash; The volume is origin of a FlexCache volume.
```

```
|flexgroup
|link:#flexgroup[flexgroup]
a|
```

```
|granular_data
|boolean
a|State of granular data on the volume. This setting is true by default when creating a new FlexGroup volume, but can be specified as false at the time of creation via a POST request. On FlexVol volumes, the setting is always false, as only FlexGroup volumes and FlexGroup constituents support this feature. Once enabled, this setting can only be disabled by restoring a Snapshot copy. Earlier versions of ONTAP (pre 9.11) 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 copy that was taken before the setting was enabled.
```

* Introduced in: 9.12

```
* x-nullable: true

|guarantee
|link:#guarantee[guarantee]
a|

|is_object_store
|boolean
a|Specifies whether the volume is provisioned for an object store server.

|is_svm_root
|boolean
a|Specifies whether the volume is a root volume of the SVM it belongs to.

|language
|string
a|Language encoding setting for volume. If no language is specified, the
volume inherits its SVM language encoding setting.

|max_dir_size
|integer
a|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
|link:#metric[metric]
a|Performance numbers, such as IOPS, latency and throughput.

|movement
|link:#movement[movement]
a|Volume movement. All attributes are modify, that is, not writable
through POST. Set PATCH state to destination_aggregate to initiate a
volume move operation. Volume movement on FlexGroup constituents are not
supported.

|msid
|integer
a|The volume's Master Set ID.
```

|name
|string
a|Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.

|nas
|link:#nas[nas]
a|

|optimize_aggregates
|boolean
a|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
|link:#qos[qos]
a|QoS information

|queue_for_encryption
|boolean
a|Specifies whether the volume is queued for encryption.

|quota
|link:#quota[quota]
a|Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

|rebalancing
|link:#rebalancing[rebalancing]
a|Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

```

|scheduled_snapshot_naming_scheme
|string
a|Naming Scheme for automatic Snapshot copies:

* create_time - Automatic Snapshot copies are saved as per the start of
their current date and time.
* ordinal - Latest automatic snapshot copy is saved as
+++<scheduled_frequency>+++0 and subsequent copies will follow the
create_time naming convention.+++</scheduled_frequency>+++

|size
|integer
a|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
|link:#snaplock[snaplock]
a|

|snapmirror
|link:#snapmirror[snapmirror]
a|Specifies attributes for SnapMirror protection.

|snapshot_count
|integer
a|Number of Snapshot copies in the volume.

|snapshot_directory_access_enabled
|boolean
a|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
a|Specifies whether or not snapshot copy locking is enabled on the volume.

|snapshot_policy

```

|link:#snapshot_policy[snapshot_policy]

a|This is a reference to the Snapshot copy policy.

|space

|link:#space[space]

a|

|state

|string

a|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

|link:#statistics[statistics]

a|These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

|status

|array[string]

a|Describes the current status of a volume.

|style

|string

a|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 constituents.

|svm
|link:#svm[svm]
a|SVM containing the volume. Required on POST.

|tiering
|link:#tiering[tiering]
a|

|type
|string
a|Type of the volume.
rw ‐ read-write volume.
dp ‐ data-protection volume.
ls ‐ load-sharing `dp` volume. Valid in GET.

|use_mirrored_aggregates
|boolean
a|Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.

|uuid
|string
a|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

|===

```

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "activity_tracking": {
    "state": "off",
    "unsupported_reason": {
      "code": "124518405",
      "message": "Volume activity tracking cannot be enabled on volumes
that contain LUNs."
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "aggressive_readahead_mode": "none",
  "analytics": {
    "initialization": {
      "state": "running"
    },
    "scan_progress": 17,
    "state": "unknown",
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes that
contain LUNs."
    }
  },
  "anti_ransomware": {

```



```
"attack_probability": "none",
"attack_reports": {
  "_links": {
    "suspects": {
      "href": "/api/resourcelink"
    }
  },
  "time": "2021-06-01 15:06:41 +0000"
},
"dry_run_start_time": "string",
"space": {
  "snapshot_count": 0,
  "used": 0,
  "used_by_logs": 0,
  "used_by_snapshots": 0
},
"state": "disabled",
"suspect_files": {
  "count": 0,
  "entropy": "string",
  "format": "string"
}
},
"anti_ransomware_state": "disabled",
"application": {
  "name": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
},
"autosize": {
  "mode": "grow"
},
"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": "svml",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"parent_volume": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
},
"split_complete_percent": 0,
"split_estimate": 0
},
"cloud_retrieval_policy": "default",
"comment": "string",
"consistency_group": {
    "name": "consistency_group_1",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
},
"constituents": {
    "aggregates": {
        "name": "string",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "movement": {
        "cutover_window": 30,
        "destination_aggregate": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": 0,
    "state": "replicating",
    "tiering_policy": "all"
},
"name": "string",
"space": {

```

```
"available": 0,
"block_storage_inactive_user_data": 0,
"capacity_tier_footprint": 0,
"footprint": 0,
"local_tier_footprint": 0,
"logical_space": {
  "available": 0,
  "used_by_afs": 0
},
"metadata": 0,
"over_provisioned": 0,
"performance_tier_footprint": 0,
"snapshot": {
  "used": 0
},
"total_footprint": 0,
"used": 0
}
},
"create_time": "2018-06-04 19:00:00 +0000",
"efficiency": {
  "application_io_size": "8k",
  "auto_state": "auto",
  "compaction": "inline",
  "compression": "inline",
  "compression_type": "none",
  "cross_volume_dedupe": "inline",
  "dedupe": "inline",
  "idcs_scanner": {
    "mode": "default",
    "operation_state": "idle",
    "status": "success",
    "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": "idle",
  "progress": "string",
  "scanner": {
    "state": "idle"
  },
  "schedule": "string",
  "space_savings": {
```

```

    "compression": 0,
    "compression_percent": 0,
    "dedupe": 0,
    "dedupe_percent": 0,
    "dedupe_sharing": 0,
    "total": 0,
    "total_percent": 0
  },
  "state": "disabled",
  "storage_efficiency_mode": "default",
  "type": "regular",
  "volume_path": "string"
},
"encryption": {
  "key_create_time": "2022-01-01 19:00:00 +0000",
  "key_id": "string",
  "key_manager_attribute": "CRN=v1:bluemix:public:containers-
kubernetes:us-south:a/asdfghjkl1234:asdfghjkl1234:worker:kubernetes-
asdfghjkl-worker1",
  "state": "encrypted",
  "status": {
    "code": "string",
    "message": "string"
  },
  "type": "none"
},
"files": {
  "used": 0
},
"flash_pool": {
  "cache_eligibility": "read",
  "cache_retention_priority": "normal",
  "caching_policy": "none"
},
"flexcache_endpoint_type": "none",
"flexgroup": {
  "name": "my_flexgroup",
  "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
},
"guarantee": {
  "type": "volume"
},
"language": "ar",
"metric": {
  "_links": {
    "self": {

```

```
    "href": "/api/resourcelink"
  }
},
"cloud": {
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"duration": "PT15S",
"flexcache": {
  "bandwidth_savings": 4096,
  "cache_miss_percent": 20,
  "duration": "PT1D",
  "status": "ok",
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"iops": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"latency": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"status": "ok",
"throughput": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 +0000"
},
"movement": {
  "cutover_window": 30,
```

```

"destination_aggregate": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "aggr1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"percent_complete": 0,
"start_time": "2020-12-07 08:45:12 +0000",
"state": "replicating",
"tiering_policy": "all"
},
"name": "vol_cs_dept",
"nas": {
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "id": 100,
  "name": "default"
},
"junction_parent": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "vs1_root",
  "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
},
"path": "/user/my_volume",
"security_style": "mixed",
"unix_permissions": 493
},
"qos": {
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "max_throughput_iops": 10000,

```

```
    "max_throughput_mbps": 500,
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"quota": {
  "state": "corrupt"
},
"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 19:00:00 +0000"
      },
      "most_recent_start_time": "2018-06-04 19:00:00 +0000"
    },
    "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,
  "notices": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "runtime": "string",
  "state": "rebalancing",
  "stop_time": "string",
  "target_used": 0,
  "used_for_imbalance": 0
},
"scheduled_snapshot_naming_scheme": "create_time",
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04 19:00:00 +0000",
  "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,
  "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
  },
  "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": "try",
    "defer_delete": "scheduled",
    "delete_order": "newest_first",
    "prefix": "string",
    "trigger": "volume"
  },
  "autodelete_trigger": "volume",
  "reserve_available": 0,
  "reserve_size": 0,
  "space_used_percent": 0,
  "used": 0
},
"snapshot_reserve_unusable": 0,
"snapshot_spill": 0,
"total_footprint": 0,
"used": 0,
"user_data": 0,
"volume_guarantee_footprint": 0
},
"state": "error",
"statistics": {
  "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,
```



```
"<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,

```



```
"<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 11:20:13 +0000"
  },
  "flexcache_raw": {
    "cache_miss_blocks": 10,
    "client_requested_blocks": 500,
    "status": "ok",
    "timestamp": "2017-01-25 11:20:13 +0000"
  },
  "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
  },
}
```



```
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,
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 11:20:13 +0000"
},
"status": {
},
"style": "flexvol",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"tiering": {
  "object_tags": {
  },
  "policy": "all",
  "storage_class": "default"
},
"type": "rw",
"uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
====

```

== Response

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====
```

=== Headers

```
[cols=3*,options=header]
|===
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|===
```

== Error

Status: Default

ONTAP Error Response Codes

```
|===
| Error Code | Description

| 787140
| One of "aggregates.uuid", "aggregates.name", or "style" must be
provided.
```

| 787141
| The specified "aggregates.name" and "aggregates.uuid" refer to different aggregates.

| 917526
| The volume name specified is a duplicate.

| 917829
| Volume autosize grow threshold must be larger than autosize shrink threshold.

| 917831
| Volume minimum autosize must be smaller than the maximum autosize.

| 917835
| Maximum allowed snapshot.reserve_percent value during a volume creation is 90. Use PATCH to set it to a higher value after the volume has been created.

| 918191
| Flexvol tiering min cooling days requires an effective cluster version of ONTAP 9.4 or later.

| 918194
| Tiering min cooling days not supported for SVMDR.

| 918195
| Tiering min cooling days not supported for non data volumes.

| 918196
| Tiering min cooling days not allowed for the provided tiering policy.

| 918215
| FlexGroup tiering min cooling days requires an effective cluster version of ONTAP 9.5 or later.

| 918233
| The target field cannot be specified for this operation.

| 918236
| The specified "parent_volume.uuid" and "parent_volume.name" do not refer to the same volume.

| 918240
| The target style is an invalid volume style.

| 918241
| The target style is an unsupported volume style for volume creation.

| 918242
| When creating a flexible volume, exactly one aggregate must be specified via either "aggregates.name" or "aggregates.uuid".

| 918243
| The specified Snapshot copy UUID is not correct for the specified Snapshot copy name.

| 918244
| Invalid "volume.type" for clone volume.

| 918246
| "volume.clone.parent_volume.name" or "volume.clone.parent_volume.uuid" must be provided.

| 918247
| Specifying a value is not valid for a volume FlexClone creation.

| 918252
| "nas.path" is invalid.

| 918290
| cloud retrieval policy requires an effective cluster version of 9.8 or later.

| 918291
| Invalid volume cloud retrieval policy for the provided tiering policy.

| 918292
| cloud retrieval policy not supported for non data volume.

| 918521
| The volume maximum autosize must be smaller than or equal to the maximum volume size.

| 918524
| Volume minimum autosize must be less than or equal to the current volume size.

| 2621706
| The specified "svm.uuid" and "svm.name" do not refer to the same SVM.

| 2621707
| No SVM was specified. Either "svm.name" or "svm.uuid" must be supplied.

| 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.

| 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.

| 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.
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]

```

`_links`

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#unsupported_reason]
```

```
[.api-collapsible-fifth-title]
```

```
unsupported_reason
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|If volume activity tracking is not supported on the volume, this field provides an appropriate error code.
```

```
|message
```

```
|string
```

```
a|If volume activity tracking is not supported on the volume, this field provides an error message detailing why this is the case.
```

```
|===
```

```
[#activity_tracking]
```

```
[.api-collapsible-fifth-title]
```

```
activity_tracking
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```



```

|Type
|Description

|state
|string
a|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
a|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
|link:#unsupported_reason[unsupported_reason]
a|

|===

[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```
|name  
|string  
a|
```

```
|uuid  
|string  
a|
```

```
|===
```

```
[#initialization]  
[.api-collapsible-fifth-title]  
initialization
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|state  
|string  
a|State of the analytics file system scan.
```

```
|===
```

```
[#unsupported_reason]  
[.api-collapsible-fifth-title]  
unsupported_reason
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|code  
|string  
a|If file system analytics is not supported on the volume, this field  
provides the error code explaining why.
```

```
|message
```

```
|string
a|If file system analytics is not supported on the volume, this field
provides the error message explaining why.
```

```
|===
```

```
[#analytics]
[.api-collapsible-fifth-title]
analytics
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|initialization
|link:#initialization[initialization]
a|
```

```
|scan_progress
```

```
|integer
```

```
a|Percentage of files in the volume that the file system analytics
initialization scan has processed. Only returned when the state is
`initializing`.
```

```
|state
```

```
|string
```

```
a|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
a|This field indicates whether or not file system analytics is supported
on the volume. If file system analytics is not supported, the reason will
be specified in the "analytics.unsupported_reason" field.
```

```
|unsupported_reason
|link:#unsupported_reason[unsupported_reason]
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|suspects
|link:#href[href]
a|
```

```
|===
```

```
[#anti_ransomware_attack_report]
[.api-collapsible-fifth-title]
anti_ransomware_attack_report
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```

|time
|string
a|Timestamp at which ransomware attack is observed.

|===

[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
|===
|Name
|Type
|Description

|snapshot_count
|integer
a|Total number of Anti-ransomware backup Snapshot copies.

|used
|integer
a|Total space in bytes used by the Anti-ransomware feature.

|used_by_logs
|integer
a|Space in bytes used by the Anti-ransomware analytics logs.

|used_by_snapshots
|integer
a|Space in bytes used by the Anti-ransomware backup Snapshot copies.

|===

[#suspect_files]
[.api-collapsible-fifth-title]
suspect_files

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|count
|integer
a|Total number of `suspect_files.format` files observed by the Anti-
ransomware analytics engine on the volume.

|entropy
|string
a|Indicates the entropy level of this file type.

|format
|string
a|File formats observed by the Anti-ransomware analytics engine on the
volume.

|===

[#anti_ransomware]
[.api-collapsible-fifth-title]
anti_ransomware

Anti-ransomware related information of the volume.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|attack_probability
|string
a|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[link:#anti_ransomware_attack_report[anti_ransomware_attack_report]]
a|

|dry_run_start_time
|string
a|Time when Anti-ransomware monitoring `state` is set to dry-run value for
starting evaluation mode.

|space
|link:#space[space]
a|

|state
|string
a|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
a|Indicates whether or not to set the surge values as historical values.

|suspect_files
|array[link:#suspect_files[suspect_files]]
a|

|===

[#application]
[.api-collapsible-fifth-title]

```

application

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

a|Name of the application to which the volume belongs. Available only when the volume is part of an application.

```
|uuid
```

```
|string
```

a|UUID of the application to which the volume belongs. Available only when the volume is part of an application.

```
|===
```

```
[#asynchronous_directory_delete]
```

```
[.api-collapsible-fifth-title]
```

```
asynchronous_directory_delete
```

Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

a|Specifies whether asynchronous directory delete from the client is enabled on the volume.

```
|trash_bin
```

```
|string
```

a|Name of the trash bin directory. If no "trash_bin" property is specified when enabling, the default trash bin name, "._ontaptrashbin", is used.


```
|===
```

```
[#autosize]  
[.api-collapsible-fifth-title]  
autosize
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|grow_threshold
```

```
|integer
```

a|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
```

a|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
```

a|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
```

a|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.

```
|shrink_threshold
|integer
a|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]
[.api-collapsible-fifth-title]
snapshot_reference
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#parent_svm]
[.api-collapsible-fifth-title]
parent_svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
```

```

|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#parent_volume]
[.api-collapsible-fifth-title]
parent_volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume.

|uuid
|string
a|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]
[.api-collapsible-fifth-title]
clone

[cols=3*,options=header]
|===
|Name
|Type
|Description

|inherited_physical_used
|integer
a|Inherited physical used from the clone's base snapshot.

|inherited_savings
|integer
a|Inherited savings from the clone's base snapshot.

|is_flexclone
|boolean
a|Specifies if this volume is a normal FlexVol or FlexClone. This field
needs to be set when creating a FlexClone. Valid in POST.

|parent_snapshot
|link:#snapshot_reference[snapshot_reference]
a|

|parent_svm
|link:#parent_svm[parent_svm]
a|

|parent_volume
|link:#parent_volume[parent_volume]
a|

|split_complete_percent
|integer
a|Percentage of FlexClone blocks split from its parent volume.

```

```
|split_estimate
|integer
a|Space required by the containing-aggregate to split the FlexClone
volume.
```

```
|split_initiated
|boolean
a|This field is set when split is executed on any FlexClone, that is when
the FlexClone volume is split from its parent FlexVol. This field needs to
be set for splitting a FlexClone form FlexVol. Valid in PATCH.
```

```
|===
```

```
[#consistency_group]
[.api-collapsible-fifth-title]
consistency_group
```

Consistency group the volume is part of.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|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
a|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]
[.api-collapsible-fifth-title]
aggregates

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the aggregate hosting the FlexGroup Constituent.

|uuid
|string
a|Unique identifier for the aggregate.

|===

[#destination_aggregate]
[.api-collapsible-fifth-title]
destination_aggregate

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid

```

```
|string
```

```
a|
```

```
|===
```

```
[#movement]
```

```
[.api-collapsible-fifth-title]
```

```
movement
```

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to `destination_aggregate` to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|cutover_window
```

```
|integer
```

a|Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.

```
|destination_aggregate
```

```
|link:#destination_aggregate[destination_aggregate]
```

```
a|Aggregate
```

```
|percent_complete
```

```
|integer
```

```
a|Completion percentage
```

```
|state
```

```
|string
```

a|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
a|Tiering policy for FabricPool
```

```
|===
```

```
[#logical_space]
[.api-collapsible-fifth-title]
logical_space
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|available
|integer
a|The amount of space available in this volume with storage efficiency
space considered used, in bytes.
```

```
|enforcement
|boolean
a|Specifies whether space accounting for operations on the volume is done
along with storage efficiency.
```

```
|reporting
|boolean
a|Specifies whether space reporting on the volume is done along with
storage efficiency.
```

```
|used_by_afs
|integer
a|The virtual space used by AFS alone (includes volume reserves) and along
with storage efficiency, in bytes.
```

```
|===
```



```

[#snapshot]
[.api-collapsible-fifth-title]
snapshot

[cols=3*,options=header]
|===
|Name
|Type
|Description

|autodelete_enabled
|boolean
a|Specifies whether Snapshot copy autodelete is currently enabled on this
volume.

|reserve_percent
|integer
a|The space that has been set aside as a reserve for Snapshot copy usage,
in percent.

|used
|integer
a|The total space used by Snapshot copies in the volume, in bytes.

|===

[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
|===
|Name
|Type
|Description

|afs_total
|integer
a|Total size of AFS, excluding snap-reserve, in bytes.

|available
|integer

```

a|The available space, in bytes.

|available_percent

|integer

a|The space available, as a percent.

|block_storage_inactive_user_data

|integer

a|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

a|Space used by capacity tier for this volume in the FabricPool aggregate, in bytes.

|footprint

|integer

a|Data used for this volume in the aggregate, in bytes.

|large_size_enabled

|boolean

a|Specifies whether the support for large volumes and large files is enabled on the volume.

|local_tier_footprint

|integer

a|Space used by the local tier for this volume in the aggregate, in bytes.

|logical_space

|link:#logical_space[logical_space]

a|

|metadata

|integer

a|Space used by the volume metadata in the aggregate, in bytes.

```
|over_provisioned
|integer
a|The amount of space not available for this volume in the aggregate, in
bytes.

|performance_tier_footprint
|integer
a|Space used by the performance tier for this volume in the FabricPool
aggregate, in bytes.

|size
|integer
a|Total provisioned size. The default size is equal to the minimum size of
20MB, in bytes.

|snapshot
|link:#snapshot[snapshot]
a|

|total_footprint
|integer
a|Data and metadata used for this volume in the aggregate, in bytes.

|used
|integer
a|The virtual space used (includes volume reserves) before storage
efficiency, in bytes.

|used_by_afs
|integer
a|The space used by Active Filesystem, in bytes.

|used_percent
|integer
a|The virtual space used (includes volume reserves) before storage
efficiency, as a percent.

|===
```

```

[#constituents]
[.api-collapsible-fifth-title]
constituents

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregates
|link:#aggregates[aggregates]
a|

|movement
|link:#movement[movement]
a|Volume movement. All attributes are modify, that is, not writable
through POST. Set PATCH state to destination_aggregate to initiate a
volume move operation. Volume movement on FlexGroup constituents are not
supported.

|name
|string
a|FlexGroup Constituents name.

|space
|link:#space[space]
a|

|===

[#idcs_scanner]
[.api-collapsible-fifth-title]
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.

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|enabled
|boolean
a|Specifies the administrative state of the inactive data compression
scanner.

|inactive_days
|integer
a|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
a|Specifies the mode of inactive data compression scanner. Valid for PATCH
and GET.

|operation_state
|string
a|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
a|Status of last inactive data compression scan on the volume.

|threshold_inactive_time
|string
a|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.+++</num>+++

|===

[#policy]

```

```
[.api-collapsible-fifth-title]
```

```
policy
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|Specifies the name of the efficiency policy.
```

```
|===
```

```
[#scanner]
```

```
[.api-collapsible-fifth-title]
```

```
scanner
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|compression
```

```
|boolean
```

```
a|Start compression if scanning old data. Valid for PATCH and GET. This option is not supported for FSX/CVO platforms.
```

```
|dedupe
```

```
|boolean
```

```
a|Start deduplication if scanning old data. Valid for PATCH and GET.
```

```
|scan_old_data
```

```
|boolean
```

```
a|Indicates whether or not to scan old data. Valid for PATCH and GET.
```

```
|state
```

```
|string
```

```
a|State of the volume efficiency scanner. Valid for PATCH and GET. Valid options for PATCH are "idle" and "active".
```

```
|===
```

```
[#space_savings]  
[.api-collapsible-fifth-title]  
space_savings
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|compression
```

```
|integer
```

```
a|Total disk space that is saved by compressing blocks on the referenced  
file system, in bytes.
```

```
|compression_percent
```

```
|integer
```

```
a|Percentage of total disk space that is saved by compressing blocks on  
the referenced file system.
```

```
|dedupe
```

```
|integer
```

```
a|Total disk space that is saved by deduplication and file cloning, in  
bytes.
```

```
|dedupe_percent
```

```
|integer
```

```
a|Percentage of total disk space that is saved by deduplication and file  
cloning.
```

```
|dedupe_sharing
```

```
|integer
```

```
a|Total disk space that is shared due to deduplication and file cloning.
```

```
|total
```

```
|integer
```

```
a|Total disk space saved in the volume due to deduplication, compression
```

and file cloning, in bytes.

```
|total_percent
|integer
a|Percentage of total disk space saved in the volume due to deduplication,
compression and file cloning.
```

```
|===
```

```
[#efficiency]
[.api-collapsible-fifth-title]
efficiency
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|application_io_size
|string
a|Block size to use by compression.
```

```
|auto_state
|string
a|Automatic deduplication schedule volume state.
auto &dash; Volumes with auto_state set to auto start post-process
deduplication automatically.
deprioritized &dash; Volumes with auto_state set to deprioritized do not
start post-process deduplication automatically.
```

```
|compaction
|string
a|The system can be enabled/disabled compaction.
inline &dash; Data will be compacted first and written to the volume.
none &dash; None
mixed &dash; Read only field for FlexGroups, where some of the constituent
volumes are compaction enabled and some are disabled.
```

```
|compression
|string
```


a|The system can be enabled/disabled compression.

inline ‐ Data will be compressed first and written to the volume.

background ‐ Data will be written to the volume and compressed later.

both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run.

none ‐ None

mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.

NOTE: that On volumes with container compression enabled, background compression refers to inactive data compression scan enabled on the volume.

|compression_type

|string

a|Compression type to use by compression. Valid for PATCH and GET.

|cross_volume_dedupe

|string

a|The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled.

inline ‐ Data will be cross volume deduped first and written to the volume.

background ‐ Data will be written to the volume and cross volume deduped later.

both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run.

none ‐ None

mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.

|dedupe

|string

a|The system can be enabled/disabled dedupe.

inline ‐ Data will be deduped first and written to the volume.

background ‐ Data will be written to the volume and deduped later.

both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run.

none ‐ None

mixed ‐ Read only field for FlexGroups, where some of the constituent

volumes are dedupe enabled and some are disabled.

|has_savings

|boolean

a|When true, indicates that the volume contains shared(deduplication, file clones) or compressed data.

|idcs_scanner

|link:#idcs_scanner[idcs_scanner]

a|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

a|Last sis operation begin timestamp.

|last_op_end

|string

a|Last sis operation end timestamp.

|last_op_err

|string

a|Last sis operation error text.

|last_op_size

|integer

a|Last sis operation size.

|last_op_state

|string

a|Last sis operation state.

|logging_enabled

|boolean

a|When true, indicates that space savings for any newly-written data are being logged.

```
|op_state
|string
a|Sis status of the volume.
```

```
|policy
|link:#policy[policy]
a|
```

```
|progress
|string
a|Sis progress of the volume.
```

```
|scanner
|link:#scanner[scanner]
a|
```

```
|schedule
|string
a|Schedule associated with volume.
```

```
|space_savings
|link:#space_savings[space_savings]
a|
```

```
|state
|string
a|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 &dash; All storage efficiency features are disabled.
mixed &dash; Read-only field for FlexGroup volumes, storage efficiency is
enabled on certain constituents and disabled on others.
On FSx for ONTAP and Cloud Volumes ONTAP &dash;
    &emsp; enabled &dash; All supported storage efficiency features for the
volume are enabled.
    &emsp; custom &dash; 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 &dash;
    &emsp; enabled &dash; 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
```

```
a|Storage efficiency mode used by volume. This parameter is supported only on AFF platform.
```

```
|type
```

```
|string
```

```
a|Sis Type of the volume.
```

```
|volume_path
```

```
|string
```

```
a|Absolute volume path of the volume.
```

```
|===
```

```
[#status]
```

```
[.api-collapsible-fifth-title]
```

```
status
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Encryption progress message code.
```

```
|message
```

```
|string
```

```
a|Encryption progress message.
```

```
|===
```

```

[#encryption]
[.api-collapsible-fifth-title]
encryption

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|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
a|Encryption key creation time of the volume.

|key_id
|string
a|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
a|Specifies an additional key manager attribute that is an identifier-
value pair, separated by '='. For example, CRN=unique-value. This
parameter is required when using the POST method and an IBM Key Lore key
manager is configured on the SVM.

|rekey
|boolean
a|If set to 'true', re-encrypts the volume with a new key. Valid in PATCH.

|state

```

```

|string
a|Volume encryption state.
encrypted &dash; The volume is completely encrypted.
encrypting &dash; Encryption operation is in progress.
partial &dash; Some constituents are encrypted and some are not.
Applicable only for FlexGroup volume.
rekeying. Encryption of volume with a new key is in progress.
unencrypted &dash; The volume is a plain-text one.

|status
|link:#status[status]
a|

|type
|string
a|Volume encryption type.
none &dash; The volume is a plain-text one.
volume &dash; The volume is encrypted with NVE (NetApp Volume Encryption).
aggregate &dash; The volume is encrypted with NAE (NetApp Aggregate
Encryption).

|===

[#error_state]
[.api-collapsible-fifth-title]
error_state

[cols=3*,options=header]
|===
|Name
|Type
|Description

|has_bad_blocks
|boolean
a|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
a|Indicates whether the file system has any inconsistencies.
true &dash; File system is inconsistent.

```

false ‐ File system in not inconsistent.

|===

[#files]

[.api-collapsible-fifth-title]

files

[cols=3*,options=header]

|===

|Name

|Type

|Description

|maximum

|integer

a|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

a|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]

[.api-collapsible-fifth-title]

flash_pool

[cols=3*,options=header]

|===

|Name

|Type

|Description

|cache_eligibility

|string

a|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

a|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

a|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]

[.api-collapsible-fifth-title]

flexgroup

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|Name of the FlexGroup volume that the constituent is part of.

|uuid

|string

a|Unique identifier for the FlexGroup volume that the constituent is part of.


```
|===
```

```
[#guarantee]  
[.api-collapsible-fifth-title]  
guarantee
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|honored
```

```
|boolean
```

```
a|Is the space guarantee of this volume honored in the aggregate?
```

```
|type
```

```
|string
```

```
a|The type of space guarantee of this volume in the aggregate.
```

```
|===
```

```
[#iops]  
[.api-collapsible-fifth-title]  
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#cloud]
[.api-collapsible-fifth-title]
cloud
```

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|duration
|string
```

a|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
|link:#iops[iops]
```

a|The rate of I/O operations observed at the storage object.

```
|latency
|link:#latency[latency]
```

a|The round trip latency in microseconds observed at the storage object.

```
|status
|string
```

a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#flexcache]
[.api-collapsible-fifth-title]
flexcache
```

Performance number for FlexCache used to measure cache effectiveness.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|bandwidth_savings
|integer
a|Bandwidth savings denoting the amount of data served locally by the
cache, in bytes.
```

```
|cache_miss_percent
|integer
a|Cache miss percentage.
```

```
|duration
|string
a|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
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```

|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

Performance numbers, such as IOPS, latency and throughput.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|cloud
|link:#cloud[cloud]
a|Performance numbers (IOPS and latency) for cloud store. These numbers
are relevant only for volumes hosted on FabricPools.

|duration
|string
a|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
|link:#flexcache[flexcache]
a|Performance number for FlexCache used to measure cache effectiveness.

```

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|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
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#movement]
[.api-collapsible-fifth-title]
movement
```

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cutover_window
|integer
a|Time window in seconds for cutover. The allowed range is between 30 to
300 seconds.

|destination_aggregate
|link:#destination_aggregate[destination_aggregate]
a|Aggregate

|percent_complete
|integer
a|Completion percentage

|start_time
|string
a|Start time of volume move.

|state
|string
a|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
a|Tiering policy for FabricPool

|===

```



```
[#export_policy]
[.api-collapsible-fifth-title]
export_policy
```

Export Policy

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
 |_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|id
```

```
|integer
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|===
```

```
[#junction_parent]
```

```
[.api-collapsible-fifth-title]
```

```
junction_parent
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
 |_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|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
a|Unique identifier for the parent volume.
```

```
|===
```

```
[#nas]
[.api-collapsible-fifth-title]
nas
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|export_policy
|link:#export_policy[export_policy]
a|Export Policy
```

```
|gid
|integer
a|The UNIX group ID of the volume. Valid in POST or PATCH.
```

```
|junction_parent
|link:#junction_parent[junction_parent]
a|
```

```
|path
|string
a|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
a|Security style associated with the volume. Valid in POST or PATCH.
mixed &dash; Mixed-style security
ntfs &dash; NTFS/WIndows-style security
unified &dash; Unified-style security, unified UNIX, NFS and CIFS
permissions
unix &dash; Unix-style security.

|uid
|integer
a|The UNIX user ID of the volume. Valid in POST or PATCH.

|unix_permissions
|integer
a|UNIX permissions to be viewed as an octal number. It consists of 4
digits derived by adding up bits 4 (read), 2 (write) and 1 (execute).
First digit selects the set user ID(4), set group ID (2) and sticky (1)
attributes. The second digit selects permission for the owner of the file;
the third selects permissions for other users in the same group; the
fourth for other users not in the group. Valid in POST or PATCH. For
security style "mixed" or "unix", the default setting is 0755 in octal
(493 in decimal) and for security style "ntfs", the default setting is
0000. In cases where only owner, group and other permissions are given (as
in 755, representing the second, third and fourth digit), first digit is
assumed to be zero.

|===

[#policy]
[.api-collapsible-fifth-title]
policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|max_throughput_iops

```

|integer

a|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

a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

|min_throughput_iops

|integer

a|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

a|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

a|The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.

|uuid

|string

a|The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

|===

[#qos]

[.api-collapsible-fifth-title]

qos

QoS information

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|policy
```

```
|link:#policy[policy]
```

```
a|
```

```
|===
```

```
[#quota]
```

```
[.api-collapsible-fifth-title]
```

```
quota
```

Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

a|This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVols or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVols or FlexGroup volumes when the quota state is "off".

```
|state
```

```
|string
```

a|Quota state of the volume

```
|===
```

```
[#last_error]
```

```
[.api-collapsible-fifth-title]
```

```
last_error
```

Error information for the last failed file move on the constituent.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|integer
```

```
a|Error code of the last file move error on the constiuent.
```

```
|destination
```

```
|integer
```

```
a|DSID of the destination constituent of the last file move error on the  
constiuent.
```

```
|file_id
```

```
|integer
```

```
a|File ID of the last file move error on the constiuent.
```

```
|time
```

```
|string
```

```
a|Time of the last file move error on the constiuent.
```

```
|===
```

```
[#movement]
```

```
[.api-collapsible-fifth-title]
```

```
movement
```

Properties on this constituent related to file movement.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|file_moves_started
```

```
|integer
a|Number of file moves started on this constituent.
```

```
|last_error
|link:#last_error[last_error]
a|Error information for the last failed file move on the constituent.
```

```
|most_recent_start_time
|string
a|Start time of the most recent file move on the constiuent.
```

```
|===
```

```
[#blocks_skipped]
[.api-collapsible-fifth-title]
blocks_skipped
```

Number of blocks skipped by the scanner on this constiuent due to various reasons.

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|efficiency_blocks
|integer
a|Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.
```

```
|efficiency_percent
|integer
a|Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.
```

```
|fast_truncate
|integer
a|Number of blocks skipped by the scanner on this constituent because fast truncate is currently running on files.
```

|footprint_invalid

|integer

a|Number of blocks skipped by the scanner on this constituent because of files with invalid space footprints.

|in_snapshot

|integer

a|Number of blocks skipped by the scanner on this constituent because of files in Snapshot copies.

|incompatible

|integer

a|Number of blocks skipped by the scanner on this constituent because of incompatible files.

|metadata

|integer

a|Number of blocks skipped by the scanner on this constituent because of metadata files.

|on_demand_destination

|integer

a|Number of blocks skipped by the scanner on this constituent because of on demand destination files.

|other

|integer

a|Number of blocks skipped by the scanner on this constituent for all other reasons.

|remote_cache

|integer

a|Number of blocks skipped by the scanner on this constituent because of remote caches.

|too_large

|integer

a|Number of blocks skipped by the scanner on this constituent because of

files that are larger than `rebalancing.max_file_size`.

|`too_small`

|integer

a|Number of blocks skipped by the scanner on this constituent because of files that are smaller than `rebalancing.min_file_size`.

|`write_fenced`

|integer

a|Number of blocks skipped by the scanner on this constituent because of files fenced for write operations.

|===

[#files_skipped]

[.api-collapsible-fifth-title]

files_skipped

Number of files skipped by the scanner on this constituent due to various reasons.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|`efficiency_blocks`

|integer

a|Number of files skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.

|`efficiency_percent`

|integer

a|Number of files skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.

|`fast_truncate`

|integer

a|Number of files skipped by the scanner on this constituent because fast

truncate is currently running on the file.

|footprint_invalid

|integer

a|Number of files skipped by the scanner on this constituent because their space footprints are invalid.

|in_snapshot

|integer

a|Number of files skipped by the scanner on this constituent because they are trapped in Snapshot copies.

|incompatible

|integer

a|Number of files skipped by the scanner on this constituent because they are incompatible.

|metadata

|integer

a|Number of files skipped by the scanner on this constituent because they metadata files.

|on_demand_destination

|integer

a|Number of files skipped by the scanner on this constituent because they are on demand destinations.

|other

|integer

a|Number of files skipped by the scanner on this constituent for all other reasons.

|remote_cache

|integer

a|Number of files skipped by the scanner on this constituent because they are remote caches.

|too_large

|integer

```
a|Number of files skipped by the scanner on this constituent because they are larger than rebalancing.max_file_size.
```

```
|too_small
```

```
|integer
```

```
a|Number of files skipped by the scanner on this constituent because they are smaller than rebalancing.min_file_size.
```

```
|write_fenced
```

```
|integer
```

```
a|Number of files skipped by the scanner on this constituent because they are fenced for write operations.
```

```
|===
```

```
[#scanner]
```

```
[.api-collapsible-fifth-title]
```

```
scanner
```

Properties related to determining which files to move and where to move them to.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|blocks_scanned
```

```
|integer
```

```
a|Number of blocks scanned on this constituent.
```

```
|blocks_skipped
```

```
|link:#blocks_skipped[blocks_skipped]
```

```
a|Number of blocks skipped by the scanner on this constituent due to various reasons.
```

```
|files_scanned
```

```
|integer
```

```
a|Number of files scanned on this constituent.
```

```
|files_skipped
|link:#files_skipped[files_skipped]
a|Number of files skipped by the scanner on this constituent due to
various reasons.
```

```
|===
```

```
[#engine]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|movement
```

```
|link:#movement[movement]
```

```
a|Properties on this constituent related to file movement.
```

```
|scanner
```

```
|link:#scanner[scanner]
```

```
a|Properties related to determining which files to move and where to move
them to.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===  
|Name  
|Type  
|Description  
  
|code  
|string  
a|Argument code
```

```
|message  
|string  
a|Message argument
```

```
|===
```

```
[#error]  
[.api-collapsible-fifth-title]  
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|arguments  
|array[link:#error_arguments[error_arguments]]  
a|Message arguments
```

```
|code  
|string  
a|Error code
```

```
|message  
|string  
a|Error message
```

```
|target  
|string  
a|The target parameter that caused the error.
```

```
|===
```

```
[#rebalancing]  
[.api-collapsible-fifth-title]  
rebalancing
```

Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|data_moved  
|integer
```

a|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  
|link:#engine[engine]
```

a|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
```

a|Specifies whether or not to exclude files that are stuck in Snapshot copies 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
a|Represents the percentage the volume is out of balance.

|imbalance_size
|integer
a|Represents how much the volume is out of balance, in bytes.

|max_constituent_imbalance_percent
|integer
a|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
a|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
a|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>+++

|max_threshold
|integer

a|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

a|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.

|min_threshold

|integer

a|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[link:#error[error]]

a|Capacity rebalancing notice messages.

|runtime

|string

a|Duration the capacity rebalancing operation has been running.

|start_time

|string

a|Time when the current capacity rebalancing operation started, or when a future scheduled rebalancing operation begins.

|state

|string

a|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.

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.

idle ‐ capacity rebalancing is running on the constituent, however, no active scanning or file movement is currently occurring.

scanning ‐ the constituent's file system is being scanned to find files to move and determine free space.

rebalancing_source ‐ a file is being moved off of the constituent.

rebalancing_dest ‐ a file is being moved to the constituent.

not_running ‐ capacity rebalancing is not running on the constituent.

unknown ‐ the system was unable to determine the rebalancing state

for the constituent.

|stop_time

|string

a|Time when the capacity rebalancing operation stopped.

|target_used

|integer

a|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

a|Represents the used size of each constituent, as determined by the rebalancing engine. Calculated by subtracting the size used by Snapshot copies, the size of files pending deletion and the size of filesystem metadata from the volume used size.

|===

[#retention]

[.api-collapsible-fifth-title]

retention

[cols=3*,options=header]

|===

|Name

|Type

|Description

|default

|string

a|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>+++

|maximum

|string

a|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>+++

|minimum

|string

a|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]  
[.api-collapsible-fifth-title]  
snaplock
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|append_mode_enabled
```

```
|boolean
```

a|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.

```
|autocommit_period
```

```
|string
```

a|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
```

a|This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.

```
|expiry_time
|string
a|Expiry time of the volume.

|is_audit_log
|boolean
a|Indicates if this volume has been configured as SnapLock audit log
volume for the SVM .

|litigation_count
|integer
a|Litigation count indicates the number of active legal-holds on the
volume.

|privileged_delete
|string
a|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
|link:#retention[retention]
a|

|type
|string
a|The SnapLock type of the volume.
compliance &dash; A SnapLock Compliance(SLC) volume provides the highest
level of WORM protection and an administrator cannot destroy a SLC volume
if it contains unexpired WORM files.
enterprise &dash; An administrator can delete a SnapLock Enterprise(SLE)
volume.
non_snaplock &dash; Indicates the volume is non-snaplock.

|unspecified_retention_file_count
|integer
a|Indicates the number of files with an unspecified retention time in the
volume.
```

```
|===
```

```
[#destinations]
```

```
[.api-collapsible-fifth-title]
```

```
destinations
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|is_cloud
```

```
|boolean
```

```
a|Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to a cloud destination.
```

```
|is_ontap
```

```
|boolean
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
snapmirror
```

```
Specifies attributes for SnapMirror protection.
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|destinations
```

```
|link:#destinations[destinations]
```

```
a|
```

```
|is_protected
```

```
|boolean
```

```
a|Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data.
```

```
|===
```

```
[#snapshot_policy]
```

```
[.api-collapsible-fifth-title]
```

```
snapshot_policy
```

This is a reference to the Snapshot copy policy.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#logical_space]
```

```
[.api-collapsible-fifth-title]
```

```
logical_space
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

Type	Description
available	
integer	a The amount of space available in this volume with storage efficiency space considered used, in bytes.
enforcement	
boolean	a Specifies whether space accounting for operations on the volume is done along with storage efficiency.
reporting	
boolean	a Specifies whether space reporting on the volume is done along with storage efficiency.
used	
integer	a SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), in bytes.
used_by_afs	
integer	a The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes.
used_by_snapshots	
integer	a Size that is logically used across all Snapshot copies in the volume, in bytes.
used_percent	
integer	a SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), as a percentage.
===	


```
[#autodelete]
[.api-collapsible-fifth-title]
autodelete
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|commitment
```

```
|string
```

a|By default, Snapshot copy autodelete does not delete Snapshot copies locked by Snapmirror, clones of a volume, a LUN, an NVMe namespace, or a file. Deletion of Snapshot copies locked by these applications is specified using this option. The default value is try.

```
|defer_delete
```

```
|string
```

a|Allows the user to inform Snapshot copy autodelete to defer the deletion of a specified Snapshot copy until the end. The default value is user_created.

```
|delete_order
```

```
|string
```

a|Specifies the order in which Snapshot copy autodelete occurs. Ordering is done using the date and time the Snapshot copy is created. The default value is oldest_first.

```
|enabled
```

```
|boolean
```

a|Specifies whether Snapshot copy autodelete is currently enabled on this volume.

```
|prefix
```

```
|string
```

a|Specifies the prefix of the Snapshot copy which if matched, is deleted last. Used with autodelete_defer_delete when used with a prefix value.

```
|target_free_space
```

```
|integer
```

a|Snapshot copies 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

a|Specifies when the system should trigger an autodelete of Snapshot copies. When set to `_volume_`, autodelete is triggered based on volume fullness. When set to `_snap_reserve_`, autodelete is triggered based on Snapshot copy reserve fullness. The default value is `_volume_`.

|===

[#snapshot]

[.api-collapsible-fifth-title]

snapshot

[cols=3*,options=header]

|===

|Name

|Type

|Description

|autodelete

|link:#autodelete[autodelete]

a|

|autodelete_enabled

|boolean

a|Specifies whether Snapshot copy 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

a|Specifies when the system should trigger an autodelete of Snapshot copies. When set to `_volume_`, autodelete is triggered based on volume fullness. When set to `_snap_reserve_`, autodelete is triggered based on Snapshot copy 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
a|Size available for Snapshot copies within the Snapshot copy reserve, in bytes.

|reserve_percent
|integer
a|The space that has been set aside as a reserve for Snapshot copy usage, in percent.

|reserve_size
|integer
a|Size in the volume that has been set aside as a reserve for Snapshot copy usage, in bytes.

|space_used_percent
|integer
a|Percentage of snapshot reserve size that has been used.

|used
|integer
a|The total space used by Snapshot copies in the volume, in bytes.

|===

[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]

|===

|Name
|Type
|Description

|afs_total
|integer
a|Total size of AFS, excluding snap-reserve, in bytes.

|auto_adaptive_compression_footprint_data_reduction

|integer

a|Savings achieved due to Auto Adaptive Compression, in bytes.

|available

|integer

a|The available space, in bytes.

|available_percent

|integer

a|The space available, as a percent.

|block_storage_inactive_user_data

|integer

a|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

a|Percentage of size that is physically used in the performance tier of the volume.

|capacity_tier_footprint

|integer

a|Space used by capacity tier for this volume in the FabricPool aggregate, in bytes.

|capacity_tier_footprint_data_reduction

|integer

a|Savings achieved in the space used by the capacity tier for this volume in the FabricPool aggregate, in bytes.

|cross_volume_dedupe_metafiles_footprint

|integer

a|Cross volume deduplication metadata footprint, in bytes.

|cross_volume_dedupe_metafiles_temporary_footprint

|integer

a|Cross volume temporary deduplication metadata footprint, in bytes.

|dedupe_metafiles_footprint

|integer

a|Deduplication metadata footprint, in bytes.

|dedupe_metafiles_temporary_footprint

|integer

a|Temporary deduplication metadata footprint, in bytes.

|delayed_free_footprint

|integer

a|Delayed free blocks footprint, in bytes.

|effective_total_footprint

|integer

a|Volume footprint after efficiency savings, in bytes.

|expected_available

|integer

a|Size that should be available for the volume, irrespective of available size in the aggregate, in bytes.

|file_operation_metadata

|integer

a|File operation metadata footprint, in bytes.

|filesystem_size

|integer

a|Total usable size of the volume, in bytes.

|filesystem_size_fixed

|boolean

a|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
a|Data used for this volume in the aggregate, in bytes.

|fractional_reserve
|integer
a|Used to change the amount of space reserved for overwrites of reserved
objects in a volume.

|full_threshold_percent
|integer
a|Volume full threshold percentage at which EMS warnings can be sent.

|is_used_stale
|boolean
a|Specifies if the virtual space used is stale.

|large_size_enabled
|boolean
a|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
a|Space used by the local tier for this volume in the aggregate, in bytes.

|logical_space
|link:#logical_space[logical_space]
a|

|metadata
|integer
a|Space used by the volume metadata in the aggregate, in bytes.

|nearly_full_threshold_percent
|integer
a|Volume nearly full threshold percentage at which EMS warnings can be
sent.
```

|over_provisioned
|integer
a|The amount of space not available for this volume in the aggregate, in bytes.

|overwrite_reserve
|integer
a|Reserved space for overwrites, in bytes.

|overwrite_reserve_used
|integer
a|Overwrite logical reserve space used, in bytes.

|percent_used
|integer
a|Percentage of the volume size that is used.

|performance_tier_footprint
|integer
a|Space used by the performance tier for this volume in the FabricPool aggregate, in bytes.

|physical_used
|integer
a|Size that is physically used in the volume, in bytes.

|physical_used_percent
|integer
a|Size that is physically used in the volume, as a percentage.

|size
|integer
a|Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes.

|size_available_for_snapshots
|integer
a|Available space for Snapshot copies from snap-reserve, in bytes.

|snapmirror_destination_footprint
|integer
a|SnapMirror destination footprint, in bytes.

|snapshot
|link:#snapshot[snapshot]
a|

|snapshot_reserve_unusable
|integer
a|Snapshot reserve that is not available for Snapshot copy creation, in bytes.

|snapshot_spill
|integer
a|Space used by the snapshot copies beyond the snap-reserve, in bytes.

|total_footprint
|integer
a|Data and metadata used for this volume in the aggregate, in bytes.

|used
|integer
a|The virtual space used (includes volume reserves) before storage efficiency, in bytes.

|used_by_afs
|integer
a|The space used by Active Filesystem, in bytes.

|user_data
|integer
a|User data, in bytes.

|volume_guarantee_footprint
|integer
a|Space reserved for future writes in the volume, in bytes.


```
|===
```

```
[#access]  
[.api-collapsible-fifth-title]  
access
```

Raw count and latency data for access operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
```

a|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]  
[.api-collapsible-fifth-title]  
audit
```

Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|count
|integer
a|Number of operations of the given type performed on this volume.
```

```
|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
dir
```

Raw count and latency data for directory-create operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|count
|integer
a|Number of operations of the given type performed on this volume.
```

```
|total_time
|integer
a|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]  
[.api-collapsible-fifth-title]  
file
```

Raw count and latency data for file-create operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
```

a|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]  
[.api-collapsible-fifth-title]  
other
```

Raw count and latency data for create operations on objects other than files, directories and symlinks.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|count
|integer
a|Number of operations of the given type performed on this volume.
```

```
|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
symlink
```

Raw count and latency data for symlink-create operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|count
|integer
a|Number of operations of the given type performed on this volume.
```

```
|total_time
|integer
a|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]  
[.api-collapsible-fifth-title]  
create
```

Raw count and latency data for create operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|dir
```

```
|link:#dir[dir]
```

a|Raw count and latency data for directory-create operations.

```
|file
```

```
|link:#file[file]
```

a|Raw count and latency data for file-create operations.

```
|other
```

```
|link:#other[other]
```

a|Raw count and latency data for create operations on objects other than files, directories and symlinks.

```
|symlink
```

```
|link:#symlink[symlink]
```

a|Raw count and latency data for symlink-create operations.

```
|===
```

```
[#getattr]  
[.api-collapsible-fifth-title]  
getattr
```

Raw count and latency data for getattr operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
link
```

```
Raw count and latency data for link operations.
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
[.api-collapsible-fifth-title]
lock
```

Raw count and latency data for lock operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|count
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
|integer
```

a|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]
[.api-collapsible-fifth-title]
lookup
```

Raw count and latency data for lookup operations.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|count
|integer
a|Number of operations of the given type performed on this volume.

|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
open

Raw count and latency data for open operations.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|Number of operations of the given type performed on this volume.

|total_time
|integer
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
read
```

Raw count and latency data for read operations, including histograms categorizing operations by size and latency.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
a|
```

```
|volume_protocol_latency_histogram_labels
```

```
|array[string]
```

```
a|Labels for the latency histogram, ranging from <2us to >20s.
```

```
|volume_protocol_size_histogram_counts
```

```
|array[integer]
```

```
a|
```

```
|volume_protocol_size_histogram_labels
```

```
|array[string]
```

```
a|Labels for the size histogram, ranging from <4KB to >1024KB.
```

```
|===
```

```
[#readdir]
```

```
[.api-collapsible-fifth-title]
```

```
readdir
```

Raw count and latency data for readdir operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
readlink
```

Raw count and latency data for readlink operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
```

a|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]
```

```
[.api-collapsible-fifth-title]
```

```
rename
```

Raw count and latency data for rename operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
```

a|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]
[.api-collapsible-fifth-title]
setattr

Raw count and latency data for setattr operations.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|count
|integer

a|Number of operations of the given type performed on this volume.

|total_time
|integer

a|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]
[.api-collapsible-fifth-title]
unlink

Raw count and latency data for unlink operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
watch
```

Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

a|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]
[.api-collapsible-fifth-title]
write

Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|count
|integer

a|Number of operations of the given type performed on this volume.

|total_time
|integer

a|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]
a|

|volume_protocol_latency_histogram_labels

```
|array[string]
a|Labels for the latency histogram, ranging from <2us to >20s.
```

```
|volume_protocol_size_histogram_counts
|array[integer]
a|
```

```
|volume_protocol_size_histogram_labels
|array[string]
a|Labels for the size histogram, ranging from <4KB to >1024KB.
```

```
|===
```

```
[#cifs_ops_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|access
```

```
|link:#access[access]
```

```
a|Raw count and latency data for access operations.
```

```
|audit
```

```
|link:#audit[audit]
```

```
a|Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.
```

```
|create
```

```
|link:#create[create]
```

```
a|Raw count and latency data for create operations.
```

```
|getattr
|link:#getattr[getattr]
a|Raw count and latency data for getattr operations.
```

```
|link
|link:#link[link]
a|Raw count and latency data for link operations.
```

```
|lock
|link:#lock[lock]
a|Raw count and latency data for lock operations.
```

```
|lookup
|link:#lookup[lookup]
a|Raw count and latency data for lookup operations.
```

```
|open
|link:#open[open]
a|Raw count and latency data for open operations.
```

```
|read
|link:#read[read]
a|Raw count and latency data for read operations, including histograms
categorizing operations by size and latency.
```

```
|readdir
|link:#readdir[readdir]
a|Raw count and latency data for readdir operations.
```

```
|readlink
|link:#readlink[readlink]
a|Raw count and latency data for readlink operations.
```

```
|rename
|link:#rename[rename]
a|Raw count and latency data for rename operations.
```



```
|setattr
|link:#setattr[setattr]
a|Raw count and latency data for setattr operations.
```

```
|unlink
|link:#unlink[unlink]
a|Raw count and latency data for unlink operations.
```

```
|watch
|link:#watch[watch]
a|Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.
```

```
|write
|link:#write[write]
a|Raw count and latency data for write operations, including histograms categorizing operations by size and latency.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#cloud]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

a|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
```

```
|link:#latency_raw[latency_raw]
```

a|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
```

a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#flexcache_raw]
[.api-collapsible-fifth-title]
flexcache_raw
```

Performance numbers for FlexCache used to measure cache effectiveness.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|cache_miss_blocks
|integer
a|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
a|Total blocks requested by the client.
```

```
|status
|string
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#nfs_ops_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|access
```

```
|link:#access[access]
```

```
a|Raw count and latency data for access operations.
```

```
|audit
```

```
|link:#audit[audit]
```

```
a|Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.
```

```
|create
```

```
|link:#create[create]
```

```
a|Raw count and latency data for create operations.
```

```
|getattr
|link:#getattr[getattr]
a|Raw count and latency data for getattr operations.
```

```
|link
|link:#link[link]
a|Raw count and latency data for link operations.
```

```
|lock
|link:#lock[lock]
a|Raw count and latency data for lock operations.
```

```
|lookup
|link:#lookup[lookup]
a|Raw count and latency data for lookup operations.
```

```
|open
|link:#open[open]
a|Raw count and latency data for open operations.
```

```
|read
|link:#read[read]
a|Raw count and latency data for read operations, including histograms
categorizing operations by size and latency.
```

```
|readdir
|link:#readdir[readdir]
a|Raw count and latency data for readdir operations.
```

```
|readlink
|link:#readlink[readlink]
a|Raw count and latency data for readlink operations.
```

```
|rename
|link:#rename[rename]
a|Raw count and latency data for rename operations.
```

```
|setattr
|link:#setattr[setattr]
a|Raw count and latency data for setattr operations.
```

```
|unlink
|link:#unlink[unlink]
a|Raw count and latency data for unlink operations.
```

```
|watch
|link:#watch[watch]
a|Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.
```

```
|write
|link:#write[write]
a|Raw count and latency data for write operations, including histograms categorizing operations by size and latency.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|cifs_ops_raw
|link:#cifs_ops_raw[cifs_ops_raw]
a|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
|link:#cloud[cloud]
a|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
|link:#flexcache_raw[flexcache_raw]
a|Performance numbers for FlexCache used to measure cache effectiveness.

|iops_raw
|link:#iops_raw[iops_raw]
a|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
|link:#latency_raw[latency_raw]
a|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
|link:#nfs_ops_raw[nfs_ops_raw]
a|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
a|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
|link:#throughput_raw[throughput_raw]
a|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
a|The timestamp of the performance data.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

SVM containing the volume. Required on POST.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#tiering]
[.api-collapsible-fifth-title]
tiering

[cols=3*,options=header]
|===

```

|Name
|Type
|Description

|min_cooling_days
|integer

a|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]

a|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".

|policy
|string

a|Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.
all ‐ This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.
auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store
none ‐ Volume blocks will not be tiered to the cloud store.
snapshot_only ‐ This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days.

```

|storage_class
|string
a|This parameter specifies the storage class that a FabricPool uses. This
feature is only available on volumes in a FabricPools on FSx or Cloud
Volumes ONTAP for AWS.

* Default value: 1
* enum: ["default", "S3_standard", "S3_standard_IA", "S3_glacier_IR"]
* Introduced in: 9.13
* x-ntap-modifyOnly: true
* x-nullable: true

|supported
|boolean
a|This parameter specifies whether or not FabricPools are selected when
provisioning a FlexGroup without specifying "aggregates.name" or
"aggregates.uuid". Only FabricPool aggregates are used if this parameter
is set to true and only non FabricPool aggregates are used if this
parameter is set to false. Tiering support for a FlexGroup can be changed
by moving all of the constituents to the required aggregates. Note that in
order to tier data, not only does the volume need to support tiering by
using FabricPools, the tiering "policy" must not be 'none'. A volume that
uses FabricPools but has a tiering "policy" of 'none' supports tiering,
but will not tier any data.

|===

[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|_tags
|array[string]
a|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

a|Indicates whether or not access time updates are enabled on the volume.

|activity_tracking

|link:#activity_tracking[activity_tracking]

a|

|aggregates

|array[link:#aggregates[aggregates]]

a|Aggregate hosting the volume. Required on POST.

|aggressive_readahead_mode

|string

a|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.

* Default value: 1

* enum: ["none", "file_prefetch"]

* Introduced in: 9.13

* x-nullable: true

|analytics

|link:#analytics[analytics]

a|

|anti_ransomware

|link:#anti_ransomware[anti_ransomware]

a|Anti-ransomware related information of the volume.

|anti_ransomware_state

|string

a|The Anti-ransomware state of the volume. If no "anti_ransomware_state" property is specified, the volume inherits the value from its parent SVM's "anti_ransomware_default_volume_state" property. If this value is "disabled", Anti-ransomware is disabled on the volume. If this value is "enabled", Anti-ransomware is enabled on the volume and alerts are raised

if any suspect is detected for those volumes. If this value is "dry_run", Anti-ransomware is enabled in the dry-run or learning mode on the volume. The "dry_run" state is same as the "enabled" state except that the analytics data is used here for learning. No alerts are raised for any detections or violations. If this value is "paused", Anti-ransomware is paused on the volume. Additionally, three more states are available, which are only valid for GET. If this value is "disable_in_progress", Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. If this value is "enable_paused", Anti-ransomware is paused on the volume from its earlier enabled state. If this value is "dry_run_paused", Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. For POST, the valid Anti-ransomware states are only "disabled", "enabled" and "dry_run", whereas for PATCH, "paused" is also valid along with the three valid states for POST.

```
|application
|link:#application[application]
a|
```

```
|asynchronous_directory_delete
|link:#asynchronous_directory_delete[asynchronous_directory_delete]
a|Configuration for asynchronous directory delete from the client. This is
only supported on Flexible volumes and FlexGroup volumes.
```

```
|autosize
|link:#autosize[autosize]
a|
```

```
|clone
|link:#clone[clone]
a|
```

```
|cloud_retrieval_policy
|string
a|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

a|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 on FSx or Cloud Volumes ONTAP.

* Introduced in: 9.13

* x-ntap-readModify: true

* x-nullable: true

|comment

|string

a|A comment for the volume. Valid in POST or PATCH.

|consistency_group

|link:#consistency_group[consistency_group]

a|Consistency group the volume is part of.

|constituents

|array[link:#constituents[constituents]]

a|FlexGroup Constituents. FlexGroup 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

a|Specifies the number of times to iterate over the aggregates listed with the "aggregates.name" or "aggregates.uuid" when creating or expanding a FlexGroup volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is

being created on multiple aggregates, the system always creates a FlexGroup volume. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list, unless 'optimize_aggregates' is specified as 'true'.

|convert_unicode

|boolean

a|Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients.

|create_time

|string

a|Creation time of the volume. This field is generated when the volume is created.

|efficiency

|link:#efficiency[efficiency]

a|

|encryption

|link:#encryption[encryption]

a|

|error_state

|link:#error_state[error_state]

a|

|files

|link:#files[files]

a|

|flash_pool

|link:#flash_pool[flash_pool]

a|

|flexcache_endpoint_type

|string

a|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

|link:#flexgroup[flexgroup]

a|

|granular_data

|boolean

a|State of granular data on the volume. This setting is true by default when creating a new FlexGroup volume, but can be specified as false at the time of creation via a POST request. On FlexVol volumes, the setting is always false, as only FlexGroup volumes and FlexGroup constituents support this feature. Once enabled, this setting can only be disabled by restoring a Snapshot copy. Earlier versions of ONTAP (pre 9.11) 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 copy that was taken before the setting was enabled.

* Introduced in: 9.12

* x-nullable: true

|guarantee

|link:#guarantee[guarantee]

a|

|is_object_store

|boolean

a|Specifies whether the volume is provisioned for an object store server.

|is_svm_root

|boolean

a|Specifies whether the volume is a root volume of the SVM it belongs to.

|language

|string

a|Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.

|max_dir_size

|integer

a|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
|link:#metric[metric]
a|Performance numbers, such as IOPS, latency and throughput.

|movement
|link:#movement[movement]
a|Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

|msid
|integer
a|The volume's Master Set ID.

|name
|string
a|Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.

|nas
|link:#nas[nas]
a|

|optimize_aggregates
|boolean
a|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
|link:#qos[qos]
a|QoS information

|queue_for_encryption

```
|boolean
a|Specifies whether the volume is queued for encryption.

|quota
|link:#quota[quota]
a|Quotas track the space or file usage of a user, group, or qtree in a
FlexVol or a FlexGroup volume.

|rebalancing
|link:#rebalancing[rebalancing]
a|Configuration and runtime properties involving non-disruptive volume
capacity rebalancing for a FlexGroup volume.

|scheduled_snapshot_naming_scheme
|string
a|Naming Scheme for automatic Snapshot copies:

* create_time - Automatic Snapshot copies are saved as per the start of
their current date and time.
* ordinal - Latest automatic snapshot copy is saved as
+++<scheduled_frequency>+++0 and subsequent copies will follow the
create_time naming convention.+++</scheduled_frequency>+++

|size
|integer
a|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
|link:#snaplock[snaplock]
a|

|snapmirror
|link:#snapmirror[snapmirror]
a|Specifies attributes for SnapMirror protection.

|snapshot_count
|integer
```

a|Number of Snapshot copies in the volume.

|snapshot_directory_access_enabled

|boolean

a|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

a|Specifies whether or not snapshot copy locking is enabled on the volume.

|snapshot_policy

|link:#snapshot_policy[snapshot_policy]

a|This is a reference to the Snapshot copy policy.

|space

|link:#space[space]

a|

|state

|string

a|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

|link:#statistics[statistics]

a|These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

|status

|array[string]

a|Describes the current status of a volume.

|style
|string
a|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 constituents.

|svm
|link:#svm[svm]
a|SVM containing the volume. Required on POST.

|tiering
|link:#tiering[tiering]
a|

|type
|string
a|Type of the volume.
rw ‐ read-write volume.
dp ‐ data-protection volume.
ls ‐ load-sharing `dp` volume. Valid in GET.

|use_mirrored_aggregates
|boolean
a|Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.

```
|uuid
|string
a|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
```

```
|===
```

```
[#job_link]
[.api-collapsible-fifth-title]
job_link
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[IDd760931d6e4340e2f47cb6bb0cb8c3ef]]
= Delete a volume
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/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.

== 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

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Unique identifier of the volume.
```

```
|force
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|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

a|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: 202, Accepted


```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
=====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code

```

```

|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
[[ID8b5350ea398936dc05fc69a210867089]]
```

```
= Retrieve a volume
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/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 [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) 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.effective_total_foorprint`
* `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`
* `tiering.storage_class`
```

== 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

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
a|Unique identifier of the volume.

|is_constituent
|boolean
|query
|False
a|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
a|Specify the fields to return.

|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|_tags
|array[string]
a|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
```


a|Indicates whether or not access time updates are enabled on the volume.

|activity_tracking
|link:#activity_tracking[activity_tracking]
a|

|aggregates
|array[link:#aggregates[aggregates]]
a|Aggregate hosting the volume. Required on POST.

|aggressive_readahead_mode
|string
a|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.

* Default value: 1
* enum: ["none", "file_prefetch"]
* Introduced in: 9.13
* x-nullable: true

|analytics
|link:#analytics[analytics]
a|

|anti_ransomware
|link:#anti_ransomware[anti_ransomware]
a|Anti-ransomware related information of the volume.

|anti_ransomware_state
|string
a|The Anti-ransomware state of the volume. If no "anti_ransomware_state" property is specified, the volume inherits the value from its parent SVM's "anti_ransomware_default_volume_state" property. If this value is "disabled", Anti-ransomware is disabled on the volume. If this value is "enabled", Anti-ransomware is enabled on the volume and alerts are raised if any suspect is detected for those volumes. If this value is "dry_run", Anti-ransomware is enabled in the dry-run or learning mode on the volume. The "dry_run" state is same as the "enabled" state except that the analytics data is used here for learning. No alerts are raised for any detections or violations. If this value is "paused", Anti-ransomware is

paused on the volume. Additionally, three more states are available, which are only valid for GET. If this value is "disable_in_progress", Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. If this value is "enable_paused", Anti-ransomware is paused on the volume from its earlier enabled state. If this value is "dry_run_paused", Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. For POST, the valid Anti-ransomware states are only "disabled", "enabled" and "dry_run", whereas for PATCH, "paused" is also valid along with the three valid states for POST.

```
|application
|link:#application[application]
a|
```

```
|asynchronous_directory_delete
|link:#asynchronous_directory_delete[asynchronous_directory_delete]
a|Configuration for asynchronous directory delete from the client. This is
only supported on Flexible volumes and FlexGroup volumes.
```

```
|autosize
|link:#autosize[autosize]
a|
```

```
|clone
|link:#clone[clone]
a|
```

```
|cloud_retrieval_policy
|string
a|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

a|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 on FSx or Cloud Volumes ONTAP.

* Introduced in: 9.13

* x-ntap-readModify: true

* x-nullable: true

|comment

|string

a|A comment for the volume. Valid in POST or PATCH.

|consistency_group

|link:#consistency_group[consistency_group]

a|Consistency group the volume is part of.

|constituents

|array[link:#constituents[constituents]]

a|FlexGroup Constituents. FlexGroup 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

a|Specifies the number of times to iterate over the aggregates listed with the "aggregates.name" or "aggregates.uuid" when creating or expanding a FlexGroup volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list, unless 'optimize_aggregates' is specified as 'true'.

```
|convert_unicode
|boolean
a|Specifies whether directory Unicode format conversion is enabled when
directories are accessed by NFS clients.

|create_time
|string
a|Creation time of the volume. This field is generated when the volume is
created.

|efficiency
|link:#efficiency[efficiency]
a|

|encryption
|link:#encryption[encryption]
a|

|error_state
|link:#error_state[error_state]
a|

|files
|link:#files[files]
a|

|flash_pool
|link:#flash_pool[flash_pool]
a|

|flexcache_endpoint_type
|string
a|FlexCache endpoint type.
none &dash; The volume is neither a FlexCache nor origin of any FlexCache.
cache &dash; The volume is a FlexCache volume.
origin &dash; The volume is origin of a FlexCache volume.

|flexgroup
|link:#flexgroup[flexgroup]
a|

|granular_data
|boolean
```

a|State of granular data on the volume. This setting is true by default when creating a new FlexGroup volume, but can be specified as false at the time of creation via a POST request. On FlexVol volumes, the setting is always false, as only FlexGroup volumes and FlexGroup constituents support this feature. Once enabled, this setting can only be disabled by restoring a Snapshot copy. Earlier versions of ONTAP (pre 9.11) 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 copy that was taken before the setting was enabled.

* Introduced in: 9.12

* x-nullable: true

|guarantee
|link:#guarantee[guarantee]
a|

|is_object_store
|boolean
a|Specifies whether the volume is provisioned for an object store server.

|is_svm_root
|boolean
a|Specifies whether the volume is a root volume of the SVM it belongs to.

|language
|string
a|Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.

|max_dir_size
|integer
a|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
|link:#metric[metric]
a|Performance numbers, such as IOPS, latency and throughput.

```
|movement
|link:#movement[movement]
a|Volume movement. All attributes are modify, that is, not writable
through POST. Set PATCH state to destination_aggregate to initiate a
volume move operation. Volume movement on FlexGroup constituents are not
supported.

|msid
|integer
a|The volume's Master Set ID.

|name
|string
a|Volume name. The name of volume must start with an alphabetic character
(a to z or A to Z) or an underscore (_). The name must be 197 or fewer
characters in length for FlexGroups, and 203 or fewer characters in length
for all other types of volumes. Volume names must be unique within an SVM.
Required on POST.

|nas
|link:#nas[nas]
a|

|optimize_aggregates
|boolean
a|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
|link:#qos[qos]
a|QoS information

|queue_for_encryption
|boolean
a|Specifies whether the volume is queued for encryption.

|quota
```

```
|link:#quota[quota]
a|Quotas track the space or file usage of a user, group, or qtree in a
FlexVol or a FlexGroup volume.

|rebalancing
|link:#rebalancing[rebalancing]
a|Configuration and runtime properties involving non-disruptive volume
capacity rebalancing for a FlexGroup volume.

|scheduled_snapshot_naming_scheme
|string
a|Naming Scheme for automatic Snapshot copies:

* create_time - Automatic Snapshot copies are saved as per the start of
their current date and time.
* ordinal - Latest automatic snapshot copy is saved as
+++<scheduled_frequency>+++0 and subsequent copies will follow the
create_time naming convention.+++</scheduled_frequency>+++

|size
|integer
a|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
|link:#snaplock[snaplock]
a|

|snapmirror
|link:#snapmirror[snapmirror]
a|Specifies attributes for SnapMirror protection.

|snapshot_count
|integer
a|Number of Snapshot copies in the volume.

|snapshot_directory_access_enabled
|boolean
```

a|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

a|Specifies whether or not snapshot copy locking is enabled on the volume.

|snapshot_policy

|link:#snapshot_policy[snapshot_policy]

a|This is a reference to the Snapshot copy policy.

|space

|link:#space[space]

a|

|state

|string

a|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

|link:#statistics[statistics]

a|These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

|status

|array[string]

a|Describes the current status of a volume.

|style

|string

a|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 constituents.

|svm
|link:#svm[svm]
a|SVM containing the volume. Required on POST.

|tiering
|link:#tiering[tiering]
a|

|type
|string
a|Type of the volume.
rw ‐ read-write volume.
dp ‐ data-protection volume.
ls ‐ load-sharing `dp` volume. Valid in GET.

|use_mirrored_aggregates
|boolean
a|Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.

|uuid
|string
a|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
```

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "activity_tracking": {
    "state": "off",
    "unsupported_reason": {
      "code": "124518405",
      "message": "Volume activity tracking cannot be enabled on volumes
that contain LUNs."
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "aggressive_readahead_mode": "none",
  "analytics": {
    "initialization": {
      "state": "running"
    }
  }
}
```

```

    },
    "scan_progress": 17,
    "state": "unknown",
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes that
contain LUNs."
    }
  },
  "anti_ransomware": {
    "attack_probability": "none",
    "attack_reports": {
      "_links": {
        "suspects": {
          "href": "/api/resourcelink"
        }
      }
    },
    "time": "2021-06-01 15:06:41 +0000"
  },
  "dry_run_start_time": "string",
  "space": {
    "snapshot_count": 0,
    "used": 0,
    "used_by_logs": 0,
    "used_by_snapshots": 0
  },
  "state": "disabled",
  "suspect_files": {
    "count": 0,
    "entropy": "string",
    "format": "string"
  }
},
"anti_ransomware_state": "disabled",
"application": {
  "name": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
},
"autosize": {
  "mode": "grow"
},
"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": "default",
"comment": "string",
"consistency_group": {
  "name": "consistency_group_1",
  "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
},
"constituents": {
  "aggregates": {
    "name": "string",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  },
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  }
}
```

```

    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "percent_complete": 0,
  "state": "replicating",
  "tiering_policy": "all"
},
"name": "string",
"space": {
  "available": 0,
  "block_storage_inactive_user_data": 0,
  "capacity_tier_footprint": 0,
  "footprint": 0,
  "local_tier_footprint": 0,
  "logical_space": {
    "available": 0,
    "used_by_afs": 0
  },
  "metadata": 0,
  "over_provisioned": 0,
  "performance_tier_footprint": 0,
  "snapshot": {
    "used": 0
  },
  "total_footprint": 0,
  "used": 0
}
},
"create_time": "2018-06-04 19:00:00 +0000",
"efficiency": {
  "application_io_size": "8k",
  "auto_state": "auto",
  "compaction": "inline",
  "compression": "inline",
  "compression_type": "none",
  "cross_volume_dedupe": "inline",
  "dedupe": "inline",
  "idcs_scanner": {
    "mode": "default",
    "operation_state": "idle",
    "status": "success",
    "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": "idle",
"progress": "string",
"scanner": {
  "state": "idle"
},
"schedule": "string",
"space_savings": {
  "compression": 0,
  "compression_percent": 0,
  "dedupe": 0,
  "dedupe_percent": 0,
  "dedupe_sharing": 0,
  "total": 0,
  "total_percent": 0
},
"state": "disabled",
"storage_efficiency_mode": "default",
"type": "regular",
"volume_path": "string"
},
"encryption": {
  "key_create_time": "2022-01-01 19:00:00 +0000",
  "key_id": "string",
  "key_manager_attribute": "CRN=v1:bluemix:public:containers-
kubernetes:us-south:a/asdfghjkl1234:asdfghjkl1234:worker:kubernetes-
asdfghjkl-worker1",
  "state": "encrypted",
  "status": {
    "code": "string",
    "message": "string"
  },
  "type": "none"
},
"files": {
  "used": 0
},
"flash_pool": {
  "cache_eligibility": "read",
  "cache_retention_priority": "normal",
  "caching_policy": "none"
},
"flexcache_endpoint_type": "none",
"flexgroup": {

```

```
"name": "my_flexgroup",
"uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
},
"guarantee": {
  "type": "volume"
},
"language": "ar",
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"cloud": {
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"duration": "PT15S",
"flexcache": {
  "bandwidth_savings": 4096,
  "cache_miss_percent": 20,
  "duration": "PT1D",
  "status": "ok",
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"iops": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"latency": {
  "read": 200,
  "total": 1000,
  "write": 100
},
```

```
"status": "ok",
"throughput": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 +0000"
},
"movement": {
  "cutover_window": 30,
  "destination_aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "percent_complete": 0,
  "start_time": "2020-12-07 08:45:12 +0000",
  "state": "replicating",
  "tiering_policy": "all"
},
"name": "vol_cs_dept",
"nas": {
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "junction_parent": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "vs1_root",
    "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
  },
  "path": "/user/my_volume",
  "security_style": "mixed",
```



```
"unix_permissions": 493
},
"qos": {
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"quota": {
  "state": "corrupt"
},
"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 19:00:00 +0000"
      },
      "most_recent_start_time": "2018-06-04 19:00:00 +0000"
    },
    "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,
"notices": {
  "arguments": {
    "code": "string",
    "message": "string"
  },
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
},
"runtime": "string",
"state": "rebalancing",
"stop_time": "string",
"target_used": 0,
"used_for_imbalance": 0
},
"scheduled_snapshot_naming_scheme": "create_time",

```

```
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04 19:00:00 +0000",
  "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,
  "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
  },
  "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": "try",
      "defer_delete": "scheduled",
      "delete_order": "newest_first",
      "prefix": "string",
      "trigger": "volume"
    },
    "autodelete_trigger": "volume",
    "reserve_available": 0,
    "reserve_size": 0,
    "space_used_percent": 0,
    "used": 0
  },
  "snapshot_reserve_unusable": 0,
  "snapshot_spill": 0,
  "total_footprint": 0,
  "used": 0,
  "user_data": 0,
  "volume_guarantee_footprint": 0
},
"state": "error",
"statistics": {
  "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,
    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,

```



```
"<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 11:20:13 +0000"
},
"flexcache_raw": {
  "cache_miss_blocks": 10,
  "client_requested_blocks": 500,

```

```
"status": "ok",
"timestamp": "2017-01-25 11:20:13 +0000"
},
"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,  
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,
```



```
"<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 11:20:13 +0000"
},
"status": {
},
"style": "flexvol",
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"tiering": {
    "object_tags": {
    },
    "policy": "all",
    "storage_class": "default"
},
"type": "rw",
"uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string

```

```

a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#unsupported_reason]
[.api-collapsible-fifth-title]
unsupported_reason

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|If volume activity tracking is not supported on the volume, this field
provides an appropriate error code.

|message
|string
a|If volume activity tracking is not supported on the volume, this field
provides an error message detailing why this is the case.

|===

```

```

[#activity_tracking]
[.api-collapsible-fifth-title]
activity_tracking

[cols=3*,options=header]
|===
|Name
|Type
|Description

|state
|string
a|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
a|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
|link:#unsupported_reason[unsupported_reason]
a|

|===

[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#initialization]
[.api-collapsible-fifth-title]
initialization

[cols=3*,options=header]
|===
|Name
|Type
|Description

|state
|string
a|State of the analytics file system scan.

|===

[#unsupported_reason]
[.api-collapsible-fifth-title]
unsupported_reason

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|code
|string
a|If file system analytics is not supported on the volume, this field
provides the error code explaining why.
```

```
|message
|string
a|If file system analytics is not supported on the volume, this field
provides the error message explaining why.
```

```
|===
```

```
[#analytics]
[.api-collapsible-fifth-title]
analytics
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|initialization
|link:#initialization[initialization]
a|
```

```
|scan_progress
|integer
a|Percentage of files in the volume that the file system analytics
initialization scan has processed. Only returned when the state is
`initializing`.
```

```
|state
|string
a|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

a|This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field.

|unsupported_reason

|link:#unsupported_reason[unsupported_reason]

a|

|===

[#_links]

[.api-collapsible-fifth-title]

_links

[cols=3*,options=header]

|===

|Name

|Type

|Description

|suspects

|link:#href[href]

a|

|===

[#anti_ransomware_attack_report]

[.api-collapsible-fifth-title]

anti_ransomware_attack_report

[cols=3*,options=header]

|===

|Name

```

|Type
|Description

|_links
|link:#_links[_links]
a|

|time
|string
a|Timestamp at which ransomware attack is observed.

|===

[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
|===
|Name
|Type
|Description

|snapshot_count
|integer
a|Total number of Anti-ransomware backup Snapshot copies.

|used
|integer
a|Total space in bytes used by the Anti-ransomware feature.

|used_by_logs
|integer
a|Space in bytes used by the Anti-ransomware analytics logs.

|used_by_snapshots
|integer
a|Space in bytes used by the Anti-ransomware backup Snapshot copies.

|===

```

```

[#suspect_files]
[.api-collapsible-fifth-title]
suspect_files

[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|Total number of `suspect_files.format` files observed by the Anti-
ransomware analytics engine on the volume.

|entropy
|string
a|Indicates the entropy level of this file type.

|format
|string
a|File formats observed by the Anti-ransomware analytics engine on the
volume.

|===

[#anti_ransomware]
[.api-collapsible-fifth-title]
anti_ransomware

Anti-ransomware related information of the volume.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|attack_probability
|string
a|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[link:#anti_ransomware_attack_report[anti_ransomware_attack_report]]
a|

|dry_run_start_time
|string
a|Time when Anti-ransomware monitoring `state` is set to dry-run value for
starting evaluation mode.

|space
|link:#space[space]
a|

|state
|string
a|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
a|Indicates whether or not to set the surge values as historical values.

|suspect_files
|array[link:#suspect_files[suspect_files]]

```

```

a|

|===

[#application]
[.api-collapsible-fifth-title]
application

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the application to which the volume belongs. Available only when
the volume is part of an application.

|uuid
|string
a|UUID of the application to which the volume belongs. Available only when
the volume is part of an application.

|===

[#asynchronous_directory_delete]
[.api-collapsible-fifth-title]
asynchronous_directory_delete

Configuration for asynchronous directory delete from the client. This is
only supported on Flexible volumes and FlexGroup volumes.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether asynchronous directory delete from the client is

```

enabled on the volume.

|trash_bin

|string

a|Name of the trash bin directory. If no "trash_bin" property is specified when enabling, the default trash bin name, "._ontaptrashbin", is used.

|===

[#autosize]

[.api-collapsible-fifth-title]

autosize

[cols=3*,options=header]

|===

|Name

|Type

|Description

|grow_threshold

|integer

a|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

a|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

a|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

a|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.

|shrink_threshold
|integer

a|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]
[.api-collapsible-fifth-title]
snapshot_reference

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links

|link:#_links[_links]

a|

|name
|string
a|

|uuid
|string
a|

|===

[#parent_svm]

```

[.api-collapsible-fifth-title]
parent_svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#parent_volume]
[.api-collapsible-fifth-title]
parent_volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume.

|uuid
|string

```


a|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]

[.api-collapsible-fifth-title]

clone

[cols=3*,options=header]

|===

|Name

|Type

|Description

|inherited_physical_used

|integer

a|Inherited physical used from the clone's base snapshot.

|inherited_savings

|integer

a|Inherited savings from the clone's base snapshot.

|is_flexclone

|boolean

a|Specifies if this volume is a normal FlexVol or FlexClone. This field needs to be set when creating a FlexClone. Valid in POST.

|parent_snapshot

|link:#snapshot_reference[snapshot_reference]

a|

|parent_svm

|link:#parent_svm[parent_svm]

a|

|parent_volume

```
|link:#parent_volume[parent_volume]
```

```
a|
```

```
|split_complete_percent
```

```
|integer
```

```
a|Percentage of FlexClone blocks split from its parent volume.
```

```
|split_estimate
```

```
|integer
```

```
a|Space required by the containing-aggregate to split the FlexClone volume.
```

```
|split_initiated
```

```
|boolean
```

```
a|This field is set when split is executed on any FlexClone, that is when the FlexClone volume is split from its parent FlexVol. This field needs to be set for splitting a FlexClone form FlexVol. Valid in PATCH.
```

```
|===
```

```
[#consistency_group]
```

```
[.api-collapsible-fifth-title]
```

```
consistency_group
```

Consistency group the volume is part of.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|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
```

a|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]

[.api-collapsible-fifth-title]

aggregates

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|Name of the aggregate hosting the FlexGroup Constituent.

|uuid

|string

a|Unique identifier for the aggregate.

|===

[#destination_aggregate]

[.api-collapsible-fifth-title]

destination_aggregate

Aggregate

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

```
a|
```

```
|name  
|string  
a|
```

```
|uuid  
|string  
a|
```

```
|===
```

```
[#movement]  
[.api-collapsible-fifth-title]  
movement
```

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|cutover_window  
|integer
```

```
a|Time window in seconds for cutover. The allowed range is between 30 to  
300 seconds.
```

```
|destination_aggregate  
|link:#destination_aggregate[destination_aggregate]  
a|Aggregate
```

```
|percent_complete  
|integer  
a|Completion percentage
```

```
|state  
|string
```

```
a|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
a|Tiering policy for FabricPool
```

```
|===
```

```
[#logical_space]
[.api-collapsible-fifth-title]
logical_space
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|available
|integer
a|The amount of space available in this volume with storage efficiency
space considered used, in bytes.
```

```
|enforcement
|boolean
a|Specifies whether space accounting for operations on the volume is done
along with storage efficiency.
```

```
|reporting
|boolean
a|Specifies whether space reporting on the volume is done along with
storage efficiency.
```

```
|used_by_afs
|integer
```

a|The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes.

|===

```
[#snapshot]
[.api-collapsible-fifth-title]
snapshot
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|autodelete_enabled
```

```
|boolean
```

a|Specifies whether Snapshot copy autodelete is currently enabled on this volume.

```
|reserve_percent
```

```
|integer
```

a|The space that has been set aside as a reserve for Snapshot copy usage, in percent.

```
|used
```

```
|integer
```

a|The total space used by Snapshot copies in the volume, in bytes.

|===

```
[#space]
[.api-collapsible-fifth-title]
space
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|afs_total
|integer
a|Total size of AFS, excluding snap-reserve, in bytes.

|available
|integer
a|The available space, in bytes.

|available_percent
|integer
a|The space available, as a percent.

|block_storage_inactive_user_data
|integer
a|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
a|Space used by capacity tier for this volume in the FabricPool aggregate,
in bytes.

|footprint
|integer
a|Data used for this volume in the aggregate, in bytes.

|large_size_enabled
|boolean
a|Specifies whether the support for large volumes and large files is
enabled on the volume.

|local_tier_footprint
|integer
a|Space used by the local tier for this volume in the aggregate, in bytes.

|logical_space
|link:#logical_space[logical_space]
```

a|

|metadata

|integer

a|Space used by the volume metadata in the aggregate, in bytes.

|over_provisioned

|integer

a|The amount of space not available for this volume in the aggregate, in bytes.

|performance_tier_footprint

|integer

a|Space used by the performance tier for this volume in the FabricPool aggregate, in bytes.

|size

|integer

a|Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes.

|snapshot

|link:#snapshot[snapshot]

a|

|total_footprint

|integer

a|Data and metadata used for this volume in the aggregate, in bytes.

|used

|integer

a|The virtual space used (includes volume reserves) before storage efficiency, in bytes.

|used_by_afs

|integer

a|The space used by Active Filesystem, in bytes.

|used_percent

|integer

a|The virtual space used (includes volume reserves) before storage efficiency, as a percent.

|===

```
[#constituents]
[.api-collapsible-fifth-title]
constituents
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|aggregates
|link:#aggregates[aggregates]
a|
```

```
|movement
|link:#movement[movement]
a|Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
```

```
|name
|string
a|FlexGroup Constituents name.
```

```
|space
|link:#space[space]
a|
```

|===

```
[#idcs_scanner]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

```
a|Specifies the administrative state of the inactive data compression scanner.
```

```
|inactive_days
```

```
|integer
```

```
a|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
```

```
a|Specifies the mode of inactive data compression scanner. Valid for PATCH and GET.
```

```
|operation_state
```

```
|string
```

```
a|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
```

```
a|Status of last inactive data compression scan on the volume.
```

```
|threshold_inactive_time
```

```
|string
```

```
a|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.+++</num>+++
```

```
|===
```

```
[#policy]  
[.api-collapsible-fifth-title]  
policy
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string  
a|Specifies the name of the efficiency policy.
```

```
|===
```

```
[#scanner]  
[.api-collapsible-fifth-title]  
scanner
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|compression  
|boolean  
a|Start compression if scanning old data. Valid for PATCH and GET. This  
option is not supported for FSX/CVO platforms.
```

```
|dedupe  
|boolean  
a|Start deduplication if scanning old data. Valid for PATCH and GET.
```

```
|scan_old_data  
|boolean
```

a|Indicates whether or not to scan old data. Valid for PATCH and GET.

|state

|string

a|State of the volume efficiency scanner. Valid for PATCH and GET. Valid options for PATCH are "idle" and "active".

|===

[#space_savings]

[.api-collapsible-fifth-title]

space_savings

[cols=3*,options=header]

|===

|Name

|Type

|Description

|compression

|integer

a|Total disk space that is saved by compressing blocks on the referenced file system, in bytes.

|compression_percent

|integer

a|Percentage of total disk space that is saved by compressing blocks on the referenced file system.

|dedupe

|integer

a|Total disk space that is saved by deduplication and file cloning, in bytes.

|dedupe_percent

|integer

a|Percentage of total disk space that is saved by deduplication and file cloning.

|dedupe_sharing

```

|integer
a|Total disk space that is shared due to deduplication and file cloning.

|total
|integer
a|Total disk space saved in the volume due to deduplication, compression
and file cloning, in bytes.

|total_percent
|integer
a|Percentage of total disk space saved in the volume due to deduplication,
compression and file cloning.

|===

[#efficiency]
[.api-collapsible-fifth-title]
efficiency

[cols=3*,options=header]
|===
|Name
|Type
|Description

|application_io_size
|string
a|Block size to use by compression.

|auto_state
|string
a|Automatic deduplication schedule volume state.
auto &dash; Volumes with auto_state set to auto start post-process
deduplication automatically.
deprioritized &dash; Volumes with auto_state set to deprioritized do not
start post-process deduplication automatically.

|compaction
|string
a|The system can be enabled/disabled compaction.
inline &dash; Data will be compacted first and written to the volume.

```

none ‐ None

mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.

|compression

|string

a|The system can be enabled/disabled compression.

inline ‐ Data will be compressed first and written to the volume.

background ‐ Data will be written to the volume and compressed later.

both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run.

none ‐ None

mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.

NOTE: that On volumes with container compression enabled, background compression refers to inactive data compression scan enabled on the volume.

|compression_type

|string

a|Compression type to use by compression. Valid for PATCH and GET.

|cross_volume_dedupe

|string

a|The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled.

inline ‐ Data will be cross volume deduped first and written to the volume.

background ‐ Data will be written to the volume and cross volume deduped later.

both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run.

none ‐ None

mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.

|dedupe

|string

a|The system can be enabled/disabled dedupe.

inline ‐ Data will be deduped first and written to the volume.
background ‐ Data will be written to the volume and deduped later.
both ‐ Inline dedupe dedupes the data and write to the volume,
background dedupe dedupes only the blocks on which inline dedupe is not
run.
none ‐ None
mixed ‐ Read only field for FlexGroups, where some of the constituent
volumes are dedupe enabled and some are disabled.

|has_savings

|boolean

a|When true, indicates that the volume contains shared(deduplication, file clones) or compressed data.

|idcs_scanner

|link:#idcs_scanner[idcs_scanner]

a|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

a|Last sis operation begin timestamp.

|last_op_end

|string

a|Last sis operation end timestamp.

|last_op_err

|string

a|Last sis operation error text.

|last_op_size

|integer

a|Last sis operation size.

|last_op_state

|string

a|Last sis operation state.

|logging_enabled

|boolean

a|When true, indicates that space savings for any newly-written data are being logged.

|op_state

|string

a|Sis status of the volume.

|policy

|link:#policy[policy]

a|

|progress

|string

a|Sis progress of the volume.

|scanner

|link:#scanner[scanner]

a|

|schedule

|string

a|Schedule associated with volume.

|space_savings

|link:#space_savings[space_savings]

a|

|state

|string

a|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

a|Storage efficiency mode used by volume. This parameter is supported only on AFF platform.

|type

|string

a|Sis Type of the volume.

|volume_path

|string

a|Absolute volume path of the volume.

|===

[#status]

[.api-collapsible-fifth-title]

status

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Encryption progress message code.

|message

```
|string  
a|Encryption progress message.
```

```
|===
```

```
[#encryption]  
[.api-collapsible-fifth-title]  
encryption
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|enabled
```

```
|boolean
```

a|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
```

a|Encryption key creation time of the volume.

```
|key_id
```

```
|string
```

a|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
```

a|Specifies an additional key manager attribute that is an identifier-value pair, separated by '='. For example, CRN=unique-value. This parameter is required when using the POST method and an IBM Key Lore key manager is configured on the SVM.

```
|rekey
|boolean
a|If set to 'true', re-encrypts the volume with a new key. Valid in PATCH.
```

```
|state
|string
a|Volume encryption state.
encrypted &dash; The volume is completely encrypted.
encrypting &dash; Encryption operation is in progress.
partial &dash; Some constituents are encrypted and some are not.
Applicable only for FlexGroup volume.
rekeying. Encryption of volume with a new key is in progress.
unencrypted &dash; The volume is a plain-text one.
```

```
|status
|link:#status[status]
a|
```

```
|type
|string
a|Volume encryption type.
none &dash; The volume is a plain-text one.
volume &dash; The volume is encrypted with NVE (NetApp Volume Encryption).
aggregate &dash; The volume is encrypted with NAE (NetApp Aggregate Encryption).
```

```
|===
```

```
[#error_state]
[.api-collapsible-fifth-title]
error_state
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|has_bad_blocks
|boolean
a|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

a|Indicates whether the file system has any inconsistencies.

true ‐ File system is inconsistent.

false ‐ File system is not inconsistent.

|===

[#files]

[.api-collapsible-fifth-title]

files

[cols=3*,options=header]

|===

|Name

|Type

|Description

|maximum

|integer

a|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

a|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]

[.api-collapsible-fifth-title]

flash_pool

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|cache_eligibility
```

```
|string
```

a|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
```

a|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
```

a|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]
```

```
[.api-collapsible-fifth-title]
```

```
flexgroup
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

a|Name of the FlexGroup volume that the constituent is part of.

|uuid

|string

a|Unique identifier for the FlexGroup volume that the constituent is part of.

|===

[#guarantee]

[.api-collapsible-fifth-title]

guarantee

[cols=3*,options=header]

|===

|Name

|Type

|Description

|honored

|boolean

a|Is the space guarantee of this volume honored in the aggregate?

|type

|string

a|The type of space guarantee of this volume in the aggregate.

|===

[#iops]

[.api-collapsible-fifth-title]

iops

The rate of I/O operations observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#cloud]
[.api-collapsible-fifth-title]
cloud
```

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|duration
|string
a|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
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
```



```
|string
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#flexcache]
[.api-collapsible-fifth-title]
flexcache
```

Performance number for FlexCache used to measure cache effectiveness.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|bandwidth_savings
```

```
|integer
```

```
a|Bandwidth savings denoting the amount of data served locally by the
cache, in bytes.
```

```
|cache_miss_percent
```

```
|integer
```

```
a|Cache miss percentage.
```

```
|duration
```

```
|string
a|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
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

Performance numbers, such as IOPS, latency and throughput.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|cloud
|link:#cloud[cloud]
a|Performance numbers (IOPS and latency) for cloud store. These numbers
are relevant only for volumes hosted on FabricPools.
```

```
|duration
|string
a|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
|link:#flexcache[flexcache]
a|Performance number for FlexCache used to measure cache effectiveness.
```

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|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
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#movement]
[.api-collapsible-fifth-title]
```

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|cutover_window
```

```
|integer
```

```
a|Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.
```

```
|destination_aggregate
```

```
|link:#destination_aggregate[destination_aggregate]
```

```
a|Aggregate
```

```
|percent_complete
```

```
|integer
```

```
a|Completion percentage
```

```
|start_time
```

```
|string
```

```
a|Start time of volume move.
```

```
|state
```

```
|string
```

```
a|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  
a|Tiering policy for FabricPool
```

```
|===
```

```
[#export_policy]  
[.api-collapsible-fifth-title]  
export_policy
```

Export Policy

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|id  
|integer  
a|
```

```
|name  
|string  
a|
```

```
|===
```

```
[#junction_parent]  
[.api-collapsible-fifth-title]  
junction_parent
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]
```

```

a|

|name
|string
a|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
a|Unique identifier for the parent volume.

|===

[#nas]
[.api-collapsible-fifth-title]
nas

[cols=3*,options=header]
|===
|Name
|Type
|Description

|export_policy
|link:#export_policy[export_policy]
a|Export Policy

|gid
|integer
a|The UNIX group ID of the volume. Valid in POST or PATCH.

|junction_parent
|link:#junction_parent[junction_parent]
a|

|path
|string
a|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

a|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

a|The UNIX user ID of the volume. Valid in POST or PATCH.

|unix_permissions

|integer

a|UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth digit), first digit is assumed to be zero.

|===

[#policy]

[.api-collapsible-fifth-title]

policy

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|max_throughput_iops

|integer

a|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

a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

|min_throughput_iops

|integer

a|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

a|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

a|The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.

|uuid

|string

a|The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

|===

```
[#qos]
[.api-collapsible-fifth-title]
qos
```

QoS information

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|policy
|link:#policy[policy]
a|

|===
```

```
[#quota]
[.api-collapsible-fifth-title]
quota
```

Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVols or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVols or FlexGroup volumes when the quota state is "off".

|state
|string
a|Quota state of the volume
```

```
|===
```

```
[#last_error]  
[.api-collapsible-fifth-title]  
last_error
```

Error information for the last failed file move on the constituent.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|code
```

```
|integer
```

a|Error code of the last file move error on the constituent.

```
|destination
```

```
|integer
```

a|DSID of the destination constituent of the last file move error on the constituent.

```
|file_id
```

```
|integer
```

a|File ID of the last file move error on the constituent.

```
|time
```

```
|string
```

a|Time of the last file move error on the constituent.

```
|===
```

```
[#movement]  
[.api-collapsible-fifth-title]  
movement
```

Properties on this constituent related to file movement.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|file_moves_started
```

```
|integer
```

```
a|Number of file moves started on this constituent.
```

```
|last_error
```

```
|link:#last_error[last_error]
```

```
a|Error information for the last failed file move on the constituent.
```

```
|most_recent_start_time
```

```
|string
```

```
a|Start time of the most recent file move on the constiuent.
```

```
|===
```

```
[#blocks_skipped]
```

```
[.api-collapsible-fifth-title]
```

```
blocks_skipped
```

Number of blocks skipped by the scanner on this constiuent due to various reasons.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|efficiency_blocks
```

```
|integer
```

```
a|Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.
```

```
|efficiency_percent
```

```
|integer
```

```
a|Number of blocks skipped by the scanner on this constituent because
```

storage efficiency lost, in percent, would be too high.

|fast_truncate

|integer

a|Number of blocks skipped by the scanner on this constituent because fast truncate is currently running on files.

|footprint_invalid

|integer

a|Number of blocks skipped by the scanner on this constituent because of files with invalid space footprints.

|in_snapshot

|integer

a|Number of blocks skipped by the scanner on this constituent because of files in Snapshot copies.

|incompatible

|integer

a|Number of blocks skipped by the scanner on this constituent because of incompatible files.

|metadata

|integer

a|Number of blocks skipped by the scanner on this constituent because of metadata files.

|on_demand_destination

|integer

a|Number of blocks skipped by the scanner on this constituent because of on demand destination files.

|other

|integer

a|Number of blocks skipped by the scanner on this constituent for all other reasons.

|remote_cache

|integer

a|Number of blocks skipped by the scanner on this constituent because of remote caches.

|too_large

|integer

a|Number of blocks skipped by the scanner on this constituent because of files that are larger than `rebalancing.max_file_size`.

|too_small

|integer

a|Number of blocks skipped by the scanner on this constituent because of files that are smaller than `rebalancing.min_file_size`.

|write_fenced

|integer

a|Number of blocks skipped by the scanner on this constituent because of files fenced for write operations.

|===

[#files_skipped]

[.api-collapsible-fifth-title]

files_skipped

Number of files skipped by the scanner on this constituent due to various reasons.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|efficiency_blocks

|integer

a|Number of files skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.

|efficiency_percent

|integer

a|Number of files skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.

|fast_truncate

|integer

a|Number of files skipped by the scanner on this constituent because fast truncate is currently running on the file.

|footprint_invalid

|integer

a|Number of files skipped by the scanner on this constituent because their space footprints are invalid.

|in_snapshot

|integer

a|Number of files skipped by the scanner on this constituent because they are trapped in Snapshot copies.

|incompatible

|integer

a|Number of files skipped by the scanner on this constituent because they are incompatible.

|metadata

|integer

a|Number of files skipped by the scanner on this constituent because they metadata files.

|on_demand_destination

|integer

a|Number of files skipped by the scanner on this constituent because they are on demand destinations.

|other

|integer

a|Number of files skipped by the scanner on this constituent for all other reasons.

|remote_cache

```
|integer
a|Number of files skipped by the scanner on this constituent because they
are remote caches.

|too_large
|integer
a|Number of files skipped by the scanner on this constituent because they
are larger than rebalancing.max_file_size.

|too_small
|integer
a|Number of files skipped by the scanner on this constituent because they
are smaller than rebalancing.min_file_size.

|write_fenced
|integer
a|Number of files skipped by the scanner on this constituent because they
are fenced for write operations.

|===

[#scanner]
[.api-collapsible-fifth-title]
scanner

Properties related to determining which files to move and where to move
them to.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|blocks_scanned
|integer
a|Number of blocks scanned on this constituent.

|blocks_skipped
|link:#blocks_skipped[blocks_skipped]
```


a|Number of blocks skipped by the scanner on this constiuent due to various reasons.

|files_scanned

|integer

a|Number of files scanned on this constituent.

|files_skipped

|link:#files_skipped[files_skipped]

a|Number of files skipped by the scanner on this constiuent due to various reasons.

|===

[#engine]

[.api-collapsible-fifth-title]

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.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|movement

|link:#movement[movement]

a|Properties on this constituent related to file movement.

|scanner

|link:#scanner[scanner]

a|Properties related to determining which files to move and where to move them to.

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

[#rebalancing]

[.api-collapsible-fifth-title]

rebalancing

Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|data_moved

|integer

a|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

|link:#engine[engine]

a|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

a|Specifies whether or not to exclude files that are stuck in Snapshot copies 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

a|Represents the percentage the volume is out of balance.

|imbalance_size

|integer

a|Represents how much the volume is out of balance, in bytes.

|max_constituent_imbalance_percent

|integer

a|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

a|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

a|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>+++

|max_threshold

|integer

a|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

a|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.

|min_threshold

|integer

a|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[link:#error[error]]

a|Capacity rebalancing notice messages.

|runtime
|string
a|Duration the capacity rebalancing operation has been running.

|start_time
|string
a|Time when the current capacity rebalancing operation started, or when a future scheduled rebalancing operation begins.

|state
|string
a|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.

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.

idle ‐ capacity rebalancing is running on the constituent, however,

no active scanning or file movement is currently occurring.
scanning ‐ the constituent's file system is being scanned to find files to move and determine free space.
rebalancing_source ‐ a file is being moved off of the constituent.
rebalancing_dest ‐ a file is being moved to the constituent.
not_running ‐ capacity rebalancing is not running on the constituent.
unknown ‐ the system was unable to determine the rebalancing state for the constituent.

|stop_time
|string
a|Time when the capacity rebalancing operation stopped.

|target_used
|integer
a|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
a|Represents the used size of each constituent, as determined by the rebalancing engine. Calculated by subtracting the size used by Snapshot copies, the size of files pending deletion and the size of filesystem metadata from the volume used size.

|===

[#retention]
[.api-collapsible-fifth-title]
retention

[cols=3*,options=header]

|===
|Name
|Type
|Description

|default
|string
a|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>+++

|maximum

|string

a|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>+++

|minimum

|string

a|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]
[.api-collapsible-fifth-title]
snaplock
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|append_mode_enabled
```

```
|boolean
```

a|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.

```
|autocommit_period
```

```
|string
```

a|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>+++++</num>+++
```

```
|compliance_clock_time
```

```
|string
```

```
a|This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
```

```
|expiry_time
```

```
|string
```

```
a|Expiry time of the volume.
```

```
|is_audit_log
```

```
|boolean
```

```
a|Indicates if this volume has been configured as SnapLock audit log volume for the SVM .
```

```
|litigation_count
```

```
|integer
```

```
a|Litigation count indicates the number of active legal-holds on the volume.
```

```
|privileged_delete
```

```
|string
```

```
a|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
```

```
|link:#retention[retention]
```

```
a|
```

```
|type
```

```
|string
```

```
a|The SnapLock type of the volume.
```

```
compliance &dash; A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files.
```

```
enterprise &dash; An administrator can delete a SnapLock Enterprise(SLE)
```

volume.

non_snaplock ‐ Indicates the volume is non-snaplock.

|unspecified_retention_file_count

|integer

a|Indicates the number of files with an unspecified retention time in the volume.

|===

[#destinations]

[.api-collapsible-fifth-title]

destinations

[cols=3*,options=header]

|===

|Name

|Type

|Description

|is_cloud

|boolean

a|Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data to a cloud destination.

|is_ontap

|boolean

a|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]

[.api-collapsible-fifth-title]

snapmirror

Specifies attributes for SnapMirror protection.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|destinations
|link:#destinations[destinations]
a|

|is_protected
|boolean
a|Specifies whether a volume is a SnapMirror source volume, using
SnapMirror to protect its data.
```

```
|===
```

```
[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy
```

This is a reference to the Snapshot copy policy.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```

[#logical_space]
[.api-collapsible-fifth-title]
logical_space

[cols=3*,options=header]
|===
|Name
|Type
|Description

|available
|integer
a|The amount of space available in this volume with storage efficiency
space considered used, in bytes.

|enforcement
|boolean
a|Specifies whether space accounting for operations on the volume is done
along with storage efficiency.

|reporting
|boolean
a|Specifies whether space reporting on the volume is done along with
storage efficiency.

|used
|integer
a|SUM of (physical-used, shared_refs, compression_saved_in_plane0,
vbn_zero, future_blk_cnt), in bytes.

|used_by_afs
|integer
a|The virtual space used by AFS alone (includes volume reserves) and along
with storage efficiency, in bytes.

|used_by_snapshots
|integer
a|Size that is logically used across all Snapshot copies in the volume, in
bytes.

```

```
|used_percent
|integer
a|SUM of (physical-used, shared_refs, compression_saved_in_plane0,
vbn_zero, future_blk_cnt), as a percentage.
```

```
|===
```

```
[#autodelete]
[.api-collapsible-fifth-title]
autodelete
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|commitment
```

```
|string
```

```
a|By default, Snapshot copy autodelete does not delete Snapshot copies
locked by Snapmirror, clones of a volume, a LUN, an NVMe namespace, or a
file. Deletion of Snapshot copies locked by these applications is
specified using this option. The default value is try.
```

```
|defer_delete
```

```
|string
```

```
a|Allows the user to inform Snapshot copy autodelete to defer the deletion
of a specified Snapshot copy until the end. The default value is
user_created.
```

```
|delete_order
```

```
|string
```

```
a|Specifies the order in which Snapshot copy autodelete occurs. Ordering
is done using the date and time the Snapshot copy is created. The default
value is oldest_first.
```

```
|enabled
```

```
|boolean
```

```
a|Specifies whether Snapshot copy autodelete is currently enabled on this
volume.
```

|prefix
|string
a|Specifies the prefix of the Snapshot copy which if matched, is deleted last. Used with `autodelete_defer_delete` when used with a prefix value.

|target_free_space
|integer
a|Snapshot copies 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
a|Specifies when the system should trigger an autodelete of Snapshot copies. When set to `_volume_`, autodelete is triggered based on volume fullness. When set to `_snap_reserve_`, autodelete is triggered based on Snapshot copy reserve fullness. The default value is `_volume_`.

|===

[#snapshot]
[.api-collapsible-fifth-title]
snapshot

[cols=3*,options=header]

|===

|Name
|Type
|Description

|autodelete
|link:#autodelete[autodelete]
a|

|autodelete_enabled
|boolean
a|Specifies whether Snapshot copy 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
a|Specifies when the system should trigger an autodelete of Snapshot
copies. When set to _volume_, autodelete is triggered based on volume
fullness. When set to _snap_reserve_, autodelete is triggered based on
Snapshot copy 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
a|Size available for Snapshot copies within the Snapshot copy reserve, in
bytes.
```

```
|reserve_percent
|integer
a|The space that has been set aside as a reserve for Snapshot copy usage,
in percent.
```

```
|reserve_size
|integer
a|Size in the volume that has been set aside as a reserve for Snapshot
copy usage, in bytes.
```

```
|space_used_percent
|integer
a|Percentage of snapshot reserve size that has been used.
```

```
|used
|integer
a|The total space used by Snapshot copies in the volume, in bytes.
```

```
|===
```

```
[#space]
[.api-collapsible-fifth-title]
space
```

```
[cols=3*,options=header]
|===
|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.

|cross_volume_dedupe_metafiles_footprint
|integer
a|Cross volume deduplication metadata footprint, in bytes.

|cross_volume_dedupe_metafiles_temporary_footprint
|integer
a|Cross volume temporary deduplication metadata footprint, in bytes.

|dedupe_metafiles_footprint
|integer
a|Deduplication metadata footprint, in bytes.

|dedupe_metafiles_temporary_footprint
|integer
a|Temporary deduplication metadata footprint, in bytes.

|delayed_free_footprint
|integer
a|Delayed free blocks footprint, in bytes.

|effective_total_footprint
|integer
a|Volume footprint after efficiency savings, in bytes.

|expected_available
|integer
a|Size that should be available for the volume, irrespective of available size in the aggregate, in bytes.

|file_operation_metadata
|integer
a|File operation metadata footprint, in bytes.

|filesystem_size
|integer
a|Total usable size of the volume, in bytes.

|filesystem_size_fixed
|boolean
a|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
a|Data used for this volume in the aggregate, in bytes.

|fractional_reserve
|integer
a|Used to change the amount of space reserved for overwrites of reserved objects in a volume.

|full_threshold_percent
|integer
a|Volume full threshold percentage at which EMS warnings can be sent.

|is_used_stale
|boolean
a|Specifies if the virtual space used is stale.

|large_size_enabled
|boolean
a|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
a|Space used by the local tier for this volume in the aggregate, in bytes.

|logical_space
|link:#logical_space[logical_space]
a|

|metadata
|integer

a|Space used by the volume metadata in the aggregate, in bytes.

|nearly_full_threshold_percent

|integer

a|Volume nearly full threshold percentage at which EMS warnings can be sent.

|over_provisioned

|integer

a|The amount of space not available for this volume in the aggregate, in bytes.

|overwrite_reserve

|integer

a|Reserved space for overwrites, in bytes.

|overwrite_reserve_used

|integer

a|Overwrite logical reserve space used, in bytes.

|percent_used

|integer

a|Percentage of the volume size that is used.

|performance_tier_footprint

|integer

a|Space used by the performance tier for this volume in the FabricPool aggregate, in bytes.

|physical_used

|integer

a|Size that is physically used in the volume, in bytes.

|physical_used_percent

|integer

a|Size that is physically used in the volume, as a percentage.

|size

```
|integer
a|Total provisioned size. The default size is equal to the minimum size of
20MB, in bytes.

|size_available_for_snapshots
|integer
a|Available space for Snapshot copies from snap-reserve, in bytes.

|snapmirror_destination_footprint
|integer
a|SnapMirror destination footprint, in bytes.

|snapshot
|link:#snapshot[snapshot]
a|

|snapshot_reserve_unusable
|integer
a|Snapshot reserve that is not available for Snapshot copy creation, in
bytes.

|snapshot_spill
|integer
a|Space used by the snapshot copies beyond the snap-reserve, in bytes.

|total_footprint
|integer
a|Data and metadata used for this volume in the aggregate, in bytes.

|used
|integer
a|The virtual space used (includes volume reserves) before storage
efficiency, in bytes.

|used_by_afs
|integer
a|The space used by Active Filesystem, in bytes.

|user_data
```

```
|integer
a|User data, in bytes.
```

```
|volume_guarantee_footprint
|integer
a|Space reserved for future writes in the volume, in bytes.
```

```
|===
```

```
[#access]
[.api-collapsible-fifth-title]
access
```

Raw count and latency data for access operations.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|count
|integer
a|Number of operations of the given type performed on this volume.
```

```
|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
audit
```

Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
dir
```

Raw count and latency data for directory-create operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
a|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]
[.api-collapsible-fifth-title]
file
```

Raw count and latency data for file-create operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
[.api-collapsible-fifth-title]
other
```


Raw count and latency data for create operations on objects other than files, directories and symlinks.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
symlink
```

Raw count and latency data for symlink-create operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```

|integer
a|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]
[.api-collapsible-fifth-title]
create

Raw count and latency data for create operations.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|dir
|link:#dir[dir]
a|Raw count and latency data for directory-create operations.

|file
|link:#file[file]
a|Raw count and latency data for file-create operations.

|other
|link:#other[other]
a|Raw count and latency data for create operations on objects other than
files, directories and symlinks.

|symlink
|link:#symlink[symlink]
a|Raw count and latency data for symlink-create operations.

|===

```

```
[#getattr]
[.api-collapsible-fifth-title]
getattr
```

Raw count and latency data for getattr operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
link
```

Raw count and latency data for link operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
a|Number of operations of the given type performed on this volume.

|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
lock
```

Raw count and latency data for lock operations.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|Number of operations of the given type performed on this volume.

|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
lookup
```

Raw count and latency data for lookup operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
```

a|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]
```

```
[.api-collapsible-fifth-title]
```

```
open
```

Raw count and latency data for open operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

|total_time

|integer

a|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]

[.api-collapsible-fifth-title]

read

Raw count and latency data for read operations, including histograms categorizing operations by size and latency.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|count

|integer

a|Number of operations of the given type performed on this volume.

|total_time

|integer

a|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]
```

```
a|
```

```
|volume_protocol_latency_histogram_labels
```

```
|array[string]
```

```
a|Labels for the latency histogram, ranging from <2us to >20s.
```

```
|volume_protocol_size_histogram_counts
```

```
|array[integer]
```

```
a|
```

```
|volume_protocol_size_histogram_labels
```

```
|array[string]
```

```
a|Labels for the size histogram, ranging from <4KB to >1024KB.
```

```
|===
```

```
[#readdir]
```

```
[.api-collapsible-fifth-title]
```

```
readdir
```

Raw count and latency data for readdir operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]  
[.api-collapsible-fifth-title]  
readlink
```

Raw count and latency data for readlink operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|count  
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time  
|integer
```

a|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]  
[.api-collapsible-fifth-title]  
rename
```

Raw count and latency data for rename operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type
```



```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
setattr
```

```
Raw count and latency data for setattr operations.
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
unlink
```

Raw count and latency data for unlink operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
```

a|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]
```

```
[.api-collapsible-fifth-title]
```

```
watch
```

Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
write
```

Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
a|
```

```
|volume_protocol_latency_histogram_labels
|array[string]
a|Labels for the latency histogram, ranging from <2us to >20s.
```

```
|volume_protocol_size_histogram_counts
|array[integer]
a|
```

```
|volume_protocol_size_histogram_labels
|array[string]
a|Labels for the size histogram, ranging from <4KB to >1024KB.
```

```
|===
```

```
[#cifs_ops_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|access
|link:#access[access]
a|Raw count and latency data for access operations.
```

```
|audit
```

```
|link:#audit[audit]
a|Raw count and latency data for audit operations. These statistics are
only applicable for CIFS protocol operations.

|create
|link:#create[create]
a|Raw count and latency data for create operations.

|getattr
|link:#getattr[getattr]
a|Raw count and latency data for getattr operations.

|link
|link:#link[link]
a|Raw count and latency data for link operations.

|lock
|link:#lock[lock]
a|Raw count and latency data for lock operations.

|lookup
|link:#lookup[lookup]
a|Raw count and latency data for lookup operations.

|open
|link:#open[open]
a|Raw count and latency data for open operations.

|read
|link:#read[read]
a|Raw count and latency data for read operations, including histograms
categorizing operations by size and latency.

|readdir
|link:#readdir[readdir]
a|Raw count and latency data for readdir operations.

|readlink
```

```

|link:#readlink[readlink]
a|Raw count and latency data for readlink operations.

|rename
|link:#rename[rename]
a|Raw count and latency data for rename operations.

|setattr
|link:#setattr[setattr]
a|Raw count and latency data for setattr operations.

|unlink
|link:#unlink[unlink]
a|Raw count and latency data for unlink operations.

|watch
|link:#watch[watch]
a|Raw count and latency data for watch operations. These statistics are
only applicable for CIFS protocol operations.

|write
|link:#write[write]
a|Raw count and latency data for write operations, including histograms
categorizing operations by size and latency.

|===

[#iops_raw]
[.api-collapsible-fifth-title]
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.

[cols=3*,options=header]
|===
|Name
|Type

```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
```

```
[.api-collapsible-fifth-title]
```

```
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#cloud]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|iops_raw
|link:#iops_raw[iops_raw]
a|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
|link:#latency_raw[latency_raw]
a|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
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#flexcache_raw]
[.api-collapsible-fifth-title]
flexcache_raw
```

Performance numbers for FlexCache used to measure cache effectiveness.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|cache_miss_blocks
|integer
```

a|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
```

a|Total blocks requested by the client.

```
|status
|string
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#nfs_ops_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|access
```

```
|link:#access[access]
```

```
a|Raw count and latency data for access operations.
```

```
|audit
```

```
|link:#audit[audit]
a|Raw count and latency data for audit operations. These statistics are
only applicable for CIFS protocol operations.

|create
|link:#create[create]
a|Raw count and latency data for create operations.

|getattr
|link:#getattr[getattr]
a|Raw count and latency data for getattr operations.

|link
|link:#link[link]
a|Raw count and latency data for link operations.

|lock
|link:#lock[lock]
a|Raw count and latency data for lock operations.

|lookup
|link:#lookup[lookup]
a|Raw count and latency data for lookup operations.

|open
|link:#open[open]
a|Raw count and latency data for open operations.

|read
|link:#read[read]
a|Raw count and latency data for read operations, including histograms
categorizing operations by size and latency.

|readdir
|link:#readdir[readdir]
a|Raw count and latency data for readdir operations.

|readlink
```

```

|link:#readlink[readlink]
a|Raw count and latency data for readlink operations.

|rename
|link:#rename[rename]
a|Raw count and latency data for rename operations.

|setattr
|link:#setattr[setattr]
a|Raw count and latency data for setattr operations.

|unlink
|link:#unlink[unlink]
a|Raw count and latency data for unlink operations.

|watch
|link:#watch[watch]
a|Raw count and latency data for watch operations. These statistics are
only applicable for CIFS protocol operations.

|write
|link:#write[write]
a|Raw count and latency data for write operations, including histograms
categorizing operations by size and latency.

|===

[#throughput_raw]
[.api-collapsible-fifth-title]
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.

[cols=3*,options=header]
|===
|Name
|Type

```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

```
[.api-collapsible-fifth-title]
```

```
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|cifs_ops_raw
```

```
|link:#cifs_ops_raw[cifs_ops_raw]
```

```
a|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

|link:#cloud[cloud]

a|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

|link:#flexcache_raw[flexcache_raw]

a|Performance numbers for FlexCache used to measure cache effectiveness.

|iops_raw

|link:#iops_raw[iops_raw]

a|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

|link:#latency_raw[latency_raw]

a|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

|link:#nfs_ops_raw[nfs_ops_raw]

a|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

a|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

|link:#throughput_raw[throughput_raw]

a|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

a|The timestamp of the performance data.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

SVM containing the volume. Required on POST.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

```
|===
```

```
[#tiering]  
[.api-collapsible-fifth-title]  
tiering
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|min_cooling_days
```

```
|integer
```

a|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]
```

a|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".

```
|policy
```

```
|string
```

a|Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH. all ‐ This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.

auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store
none ‐ Volume blocks will not be tiered to the cloud store.
snapshot_only ‐ This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days.

|storage_class

|string

a|This parameter specifies the storage class that a FabricPool uses. This feature is only available on volumes in a FabricPools on FSx or Cloud Volumes ONTAP for AWS.

* Default value: 1

* enum: ["default", "S3_standard", "S3_standard_IA", "S3_glacier_IR"]

* Introduced in: 9.13

* x-ntap-modifyOnly: true

* x-nullable: true

|supported

|boolean

a|This parameter specifies whether or not FabricPools are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only FabricPool aggregates are used if this parameter is set to true and only non FabricPool aggregates are used if this parameter is set to false. Tiering support for a FlexGroup can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only does the volume need to support tiering by using FabricPools, the tiering "policy" must not be 'none'. A volume that uses FabricPools but has a tiering "policy" of 'none' supports tiering, but will not tier any data.

|===

//end collapsible .Definitions block

====

[[IDb45eb64a26c41f4b74c92e01cf3d46]]

= Update volume attributes

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-block]#`/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.

== Optional properties

* `queue_for_encryption` - Queue volumes for encryption when `encryption.enabled=true`. If this option is not provided or is false, conversion of volumes starts immediately. When there are volumes in the queue and less than four encryptions are running, volumes are encrypted in the order in which they are queued.

* `encryption.action` - You can pause an ongoing rekey/conversion operation or resume a paused rekey/conversion operation using this field. The following actions are supported for this field: ‐ conversion_pause - Pause an encryption conversion operation currently in progress ‐ conversion_resume - Resume a paused encryption conversion operation ‐ rekey_pause - Pause an encryption rekey operation currently in progress ‐ rekey_resume - Resume a paused encryption rekey operation

== Related ONTAP commands

- * `volume unmount`
- * `volume mount`
- * `volume online`
- * `volume offline`
- * `volume modify`
- * `volume clone modify`
- * `volume efficiency modify`
- * `volume quota on`
- * `volume quota off`
- * `volume snaplock modify`

```
* `volume encryption conversion start`  
* `volume encryption rekey start`  
* `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

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Unique identifier of the volume.
```

```
|restore_to.snapshot.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|UUID of the Snapshot copy to restore volume to the point in time the Snapshot copy was taken.
```

```
|restore_to.snapshot.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Name of the Snapshot copy to restore volume to the point in time the Snapshot copy was taken.
```

|restore_to.path
|string
|query
|False
a|Path to the file which is restored from the Snapshot copy.

* Introduced in: 9.10

|restore_to.restore_path
|string
|query
|False
a|Specifies the destination location inside the volume where the file is restored.

* Introduced in: 9.10

|restore_to.start_byte
|number
|query
|False
a|Starting byte offset of the source file, in multiples of 4096.

* Introduced in: 9.10

|restore_to.byte_count
|number
|query
|False
a|Number of bytes to restore from the source file, in multiples of 4096.

* Introduced in: 9.10

|preserve_lun_ids
|boolean
|query
|False
a|Specifies whether LUN IDs need to be preserved during a Snapshot copy restore operation.

* Introduced in: 9.10

* Default value:

```
|nvfail
|string
|query
|False
```

a|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
```

a|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
```

```
|sizing_method
|string
|query
|False
```

a|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"]
```

```
|scheduled_snapshot_naming_scheme
|string
|query
|False
```

a|Naming Scheme for automatic Snapshot copies:

* create_time - Automatic Snapshot copies are saved as per the start of their current date and time.

* ordinal - Latest automatic snapshot copy is saved as
+++<scheduled_frequency>+++0 and subsequent copies will follow the
create_time naming convention.+++</scheduled_frequency>+++

* Introduced in: 9.10

* Default value: 1

* enum: ["create_time", "ordinal"]

</scheduled_frequency>

```
|clone.match_parent_storage_tier
|boolean
|query
|False
```

a|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

```
|aggressive_readahead_mode
|string
|query
|False
```

a|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.

* Introduced in: 9.13

* Default value: 1

* enum: ["none", "file_prefetch"]

```
|return_timeout
```

```

|integer
|query
|False
a|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

|validate_only
|boolean
|query
|False
a|Validate the operation and its parameters, without actually performing
the operation.

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|_tags
|array[string]
a|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

```

a|Indicates whether or not access time updates are enabled on the volume.

|activity_tracking
|link:#activity_tracking[activity_tracking]
a|

|aggregates
|array[link:#aggregates[aggregates]]
a|Aggregate hosting the volume. Required on POST.

|aggressive_readahead_mode
|string
a|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.

* Default value: 1
* enum: ["none", "file_prefetch"]
* Introduced in: 9.13
* x-nullable: true

|analytics
|link:#analytics[analytics]
a|

|anti_ransomware
|link:#anti_ransomware[anti_ransomware]
a|Anti-ransomware related information of the volume.

|anti_ransomware_state
|string
a|The Anti-ransomware state of the volume. If no "anti_ransomware_state" property is specified, the volume inherits the value from its parent SVM's "anti_ransomware_default_volume_state" property. If this value is "disabled", Anti-ransomware is disabled on the volume. If this value is "enabled", Anti-ransomware is enabled on the volume and alerts are raised if any suspect is detected for those volumes. If this value is "dry_run", Anti-ransomware is enabled in the dry-run or learning mode on the volume. The "dry_run" state is same as the "enabled" state except that the analytics data is used here for learning. No alerts are raised for any detections or violations. If this value is "paused", Anti-ransomware is

paused on the volume. Additionally, three more states are available, which are only valid for GET. If this value is "disable_in_progress", Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. If this value is "enable_paused", Anti-ransomware is paused on the volume from its earlier enabled state. If this value is "dry_run_paused", Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. For POST, the valid Anti-ransomware states are only "disabled", "enabled" and "dry_run", whereas for PATCH, "paused" is also valid along with the three valid states for POST.

```
|application
|link:#application[application]
a|
```

```
|asynchronous_directory_delete
|link:#asynchronous_directory_delete[asynchronous_directory_delete]
a|Configuration for asynchronous directory delete from the client. This is
only supported on Flexible volumes and FlexGroup volumes.
```

```
|autosize
|link:#autosize[autosize]
a|
```

```
|clone
|link:#clone[clone]
a|
```

```
|cloud_retrieval_policy
|string
a|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

a|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 on FSx or Cloud Volumes ONTAP.

* Introduced in: 9.13

* x-ntap-readModify: true

* x-nullable: true

|comment

|string

a|A comment for the volume. Valid in POST or PATCH.

|consistency_group

|link:#consistency_group[consistency_group]

a|Consistency group the volume is part of.

|constituents

|array[link:#constituents[constituents]]

a|FlexGroup Constituents. FlexGroup 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

a|Specifies the number of times to iterate over the aggregates listed with the "aggregates.name" or "aggregates.uuid" when creating or expanding a FlexGroup volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list, unless 'optimize_aggregates' is specified as 'true'.

```
|convert_unicode
|boolean
a|Specifies whether directory Unicode format conversion is enabled when
directories are accessed by NFS clients.

|create_time
|string
a|Creation time of the volume. This field is generated when the volume is
created.

|efficiency
|link:#efficiency[efficiency]
a|

|encryption
|link:#encryption[encryption]
a|

|error_state
|link:#error_state[error_state]
a|

|files
|link:#files[files]
a|

|flash_pool
|link:#flash_pool[flash_pool]
a|

|flexcache_endpoint_type
|string
a|FlexCache endpoint type.
none &dash; The volume is neither a FlexCache nor origin of any FlexCache.
cache &dash; The volume is a FlexCache volume.
origin &dash; The volume is origin of a FlexCache volume.

|flexgroup
|link:#flexgroup[flexgroup]
a|

|granular_data
|boolean
```

a|State of granular data on the volume. This setting is true by default when creating a new FlexGroup volume, but can be specified as false at the time of creation via a POST request. On FlexVol volumes, the setting is always false, as only FlexGroup volumes and FlexGroup constituents support this feature. Once enabled, this setting can only be disabled by restoring a Snapshot copy. Earlier versions of ONTAP (pre 9.11) 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 copy that was taken before the setting was enabled.

* Introduced in: 9.12

* x-nullable: true

|guarantee
|link:#guarantee[guarantee]
a|

|is_object_store
|boolean
a|Specifies whether the volume is provisioned for an object store server.

|is_svm_root
|boolean
a|Specifies whether the volume is a root volume of the SVM it belongs to.

|language
|string
a|Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.

|max_dir_size
|integer
a|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
|link:#metric[metric]
a|Performance numbers, such as IOPS, latency and throughput.

```
|movement
|link:#movement[movement]
a|Volume movement. All attributes are modify, that is, not writable
through POST. Set PATCH state to destination_aggregate to initiate a
volume move operation. Volume movement on FlexGroup constituents are not
supported.

|msid
|integer
a|The volume's Master Set ID.

|name
|string
a|Volume name. The name of volume must start with an alphabetic character
(a to z or A to Z) or an underscore (_). The name must be 197 or fewer
characters in length for FlexGroups, and 203 or fewer characters in length
for all other types of volumes. Volume names must be unique within an SVM.
Required on POST.

|nas
|link:#nas[nas]
a|

|optimize_aggregates
|boolean
a|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
|link:#qos[qos]
a|QoS information

|queue_for_encryption
|boolean
a|Specifies whether the volume is queued for encryption.

|quota
```

```
|link:#quota[quota]
a|Quotas track the space or file usage of a user, group, or qtree in a
FlexVol or a FlexGroup volume.

|rebalancing
|link:#rebalancing[rebalancing]
a|Configuration and runtime properties involving non-disruptive volume
capacity rebalancing for a FlexGroup volume.

|scheduled_snapshot_naming_scheme
|string
a|Naming Scheme for automatic Snapshot copies:

* create_time - Automatic Snapshot copies are saved as per the start of
their current date and time.
* ordinal - Latest automatic snapshot copy is saved as
+++<scheduled_frequency>+++0 and subsequent copies will follow the
create_time naming convention.+++</scheduled_frequency>+++

|size
|integer
a|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
|link:#snaplock[snaplock]
a|

|snapmirror
|link:#snapmirror[snapmirror]
a|Specifies attributes for SnapMirror protection.

|snapshot_count
|integer
a|Number of Snapshot copies in the volume.

|snapshot_directory_access_enabled
|boolean
```

a|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

a|Specifies whether or not snapshot copy locking is enabled on the volume.

|snapshot_policy

|link:#snapshot_policy[snapshot_policy]

a|This is a reference to the Snapshot copy policy.

|space

|link:#space[space]

a|

|state

|string

a|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

|link:#statistics[statistics]

a|These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

|status

|array[string]

a|Describes the current status of a volume.

|style

|string

a|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 constituents.

|svm
|link:#svm[svm]
a|SVM containing the volume. Required on POST.

|tiering
|link:#tiering[tiering]
a|

|type
|string
a|Type of the volume.
rw ‐ read-write volume.
dp ‐ data-protection volume.
ls ‐ load-sharing `dp` volume. Valid in GET.

|use_mirrored_aggregates
|boolean
a|Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.

|uuid
|string
a|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
```

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "_tags": [
    "team:csi",
    "environment:test"
  ],
  "activity_tracking": {
    "state": "off",
    "unsupported_reason": {
      "code": "124518405",
      "message": "Volume activity tracking cannot be enabled on volumes
that contain LUNs."
    }
  },
  "aggregates": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "aggressive_readahead_mode": "none",
  "analytics": {
    "initialization": {
      "state": "running"
    }
  }
}
```

```

    },
    "scan_progress": 17,
    "state": "unknown",
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes that
contain LUNs."
    }
  },
  "anti_ransomware": {
    "attack_probability": "none",
    "attack_reports": {
      "_links": {
        "suspects": {
          "href": "/api/resourcelink"
        }
      }
    },
    "time": "2021-06-01 15:06:41 +0000"
  },
  "dry_run_start_time": "string",
  "space": {
    "snapshot_count": 0,
    "used": 0,
    "used_by_logs": 0,
    "used_by_snapshots": 0
  },
  "state": "disabled",
  "suspect_files": {
    "count": 0,
    "entropy": "string",
    "format": "string"
  }
},
"anti_ransomware_state": "disabled",
"application": {
  "name": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
},
"autosize": {
  "mode": "grow"
},
"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": "default",
"comment": "string",
"consistency_group": {
  "name": "consistency_group_1",
  "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
},
"constituents": {
  "aggregates": {
    "name": "string",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  },
  "movement": {
    "cutover_window": 30,
    "destination_aggregate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  }
}

```

```

    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "percent_complete": 0,
  "state": "replicating",
  "tiering_policy": "all"
},
"name": "string",
"space": {
  "available": 0,
  "block_storage_inactive_user_data": 0,
  "capacity_tier_footprint": 0,
  "footprint": 0,
  "local_tier_footprint": 0,
  "logical_space": {
    "available": 0,
    "used_by_afs": 0
  },
  "metadata": 0,
  "over_provisioned": 0,
  "performance_tier_footprint": 0,
  "snapshot": {
    "used": 0
  },
  "total_footprint": 0,
  "used": 0
}
},
"create_time": "2018-06-04 19:00:00 +0000",
"efficiency": {
  "application_io_size": "8k",
  "auto_state": "auto",
  "compaction": "inline",
  "compression": "inline",
  "compression_type": "none",
  "cross_volume_dedupe": "inline",
  "dedupe": "inline",
  "idcs_scanner": {
    "mode": "default",
    "operation_state": "idle",
    "status": "success",
    "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": "idle",
"progress": "string",
"scanner": {
  "state": "idle"
},
"schedule": "string",
"space_savings": {
  "compression": 0,
  "compression_percent": 0,
  "dedupe": 0,
  "dedupe_percent": 0,
  "dedupe_sharing": 0,
  "total": 0,
  "total_percent": 0
},
"state": "disabled",
"storage_efficiency_mode": "default",
"type": "regular",
"volume_path": "string"
},
"encryption": {
  "key_create_time": "2022-01-01 19:00:00 +0000",
  "key_id": "string",
  "key_manager_attribute": "CRN=v1:bluemix:public:containers-
kubernetes:us-south:a/asdfghjkl1234:asdfghjkl1234:worker:kubernetes-
asdfghjkl-worker1",
  "state": "encrypted",
  "status": {
    "code": "string",
    "message": "string"
  },
  "type": "none"
},
"files": {
  "used": 0
},
"flash_pool": {
  "cache_eligibility": "read",
  "cache_retention_priority": "normal",
  "caching_policy": "none"
},
"flexcache_endpoint_type": "none",
"flexgroup": {

```

```
"name": "my_flexgroup",
"uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
},
"guarantee": {
  "type": "volume"
},
"language": "ar",
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"cloud": {
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"duration": "PT15S",
"flexcache": {
  "bandwidth_savings": 4096,
  "cache_miss_percent": 20,
  "duration": "PT1D",
  "status": "ok",
  "timestamp": "2017-01-25 11:20:13 +0000"
},
"iops": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"latency": {
  "read": 200,
  "total": 1000,
  "write": 100
},
```

```
"status": "ok",
"throughput": {
  "read": 200,
  "total": 1000,
  "write": 100
},
"timestamp": "2017-01-25 11:20:13 +0000"
},
"movement": {
  "cutover_window": 30,
  "destination_aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "percent_complete": 0,
  "start_time": "2020-12-07 08:45:12 +0000",
  "state": "replicating",
  "tiering_policy": "all"
},
"name": "vol_cs_dept",
"nas": {
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": 100,
    "name": "default"
  },
  "junction_parent": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "vs1_root",
    "uuid": "75c9cfb0-3eb4-11eb-9fb4-005056bb088a"
  },
  "path": "/user/my_volume",
  "security_style": "mixed",
```

```
"unix_permissions": 493
},
"qos": {
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput_iops": 10000,
    "max_throughput_mbps": 500,
    "min_throughput_iops": 2000,
    "min_throughput_mbps": 500,
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"quota": {
  "state": "corrupt"
},
"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 19:00:00 +0000"
    },
    "most_recent_start_time": "2018-06-04 19:00:00 +0000"
  },
  "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,
"notices": {
  "arguments": {
    "code": "string",
    "message": "string"
  },
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
},
"runtime": "string",
"state": "rebalancing",
"stop_time": "string",
"target_used": 0,
"used_for_imbalance": 0
},
"scheduled_snapshot_naming_scheme": "create_time",

```

```
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04 19:00:00 +0000",
  "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,
  "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
  },
  "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": "try",
      "defer_delete": "scheduled",
      "delete_order": "newest_first",
      "prefix": "string",
      "trigger": "volume"
    },
    "autodelete_trigger": "volume",
    "reserve_available": 0,
    "reserve_size": 0,
    "space_used_percent": 0,
    "used": 0
  },
  "snapshot_reserve_unusable": 0,
  "snapshot_spill": 0,
  "total_footprint": 0,
  "used": 0,
  "user_data": 0,
  "volume_guarantee_footprint": 0
},
"state": "error",
"statistics": {
  "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,
```



```
"<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,

```



```
"<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 11:20:13 +0000"
},
"flexcache_raw": {
  "cache_miss_blocks": 10,
  "client_requested_blocks": 500,

```

```
"status": "ok",
"timestamp": "2017-01-25 11:20:13 +0000"
},
"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,  
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,

```



```
"<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 11:20:13 +0000"
},
"status": {
},
"style": "flexvol",
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"tiering": {
    "object_tags": {
    },
    "policy": "all",
    "storage_class": "default"
},
"type": "rw",
"uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
====

== Response
```

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 787141
| The specified "aggregates.name" and "aggregates.uuid" refer to different aggregates.

| 787144
| The volume is on an aggregate that is not part of a FabricPool.

```

| 917505
| SVM not found.

| 917829
| Volume autosize grow threshold must be larger than autosize shrink threshold.

| 917831
| Volume minimum autosize must be smaller than the maximum autosize.

| 918193
| Cannot modify tiering min cooling days when vol move is in progress.

| 918194
| Tiering min cooling days not supported for SVMDR.

| 918195
| Tiering min cooling days not supported for non data volumes.

| 918196
| Tiering min cooling days not allowed for the provided tiering policy.

| 918248
| Specifying a value is not valid for initiating volume FlexClone split operation.

| 918251
| Specifying a value is not valid for a Snapshot copy restore operation.

| 918252
| specified "nas.path" is invalid.

| 918265
| Volume is on the same aggregate.

| 918266
| "movement.destination_aggregate" and "movement.state" are mutually exclusive, unless the state is "cutover-wait".

| 918267
| The specified "movement.destination_aggregate" does not exist.

| 918291
| Invalid volume cloud retrieval policy for the provided tiering policy.

| 918292
| cloud retrieval policy not supported for non data volume.

| 918293
| Cannot modify cloud retrieval policy when vol move is in progress.

| 918521
| The volume maximum autosize must be smaller than or equal to the maximum volume size.

| 918532
| The FlexClone match-parent-storage-tier option requires an effective cluster version of 9.9.1 or later.

| 918533
| The FlexClone match-parent-storage-tier option not applicable for FlexClone volumes hosted on non-FabricPool storage.

| 918534
| The tiering policy values are different for the FlexClone volume and its parent volume. The match-parent-storage-tier option cannot be set to true.

| 918535
| The tiering minimum cooling day values are different for the FlexClone volume and its parent volume. The match-parent-storage-tier option cannot be set to true.

| 918537
| Could not get the FlexClone volume tiering policy or its parent volume tiering policy. Wait a minute and try again.

| 918538
| The match-parent-storage-tier option is not supported for clone creation.

| 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\" or \"archive\" 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.

| 2424998
| Unable to determine whether MetroCluster is configured.

| 9437885
| The volume is not online.

| 13107256
| Operation is only supported on 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 copy 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.

| 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 copy.

| 13109260
| Failed to enable granular data on the volume.

| 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.

| 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.

| 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.

| 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.
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```
.Example error
```

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{  
  "error": {  
    "arguments": {  
      "code": "string",  
      "message": "string"  
    },  
    "code": "4",  
    "message": "entry doesn't exist",  
    "target": "uuid"  
  }  
}
```

```
=====
```

```
== Definitions
```

```
[.api-def-first-level]
```

```
.See Definitions
```

```
[%collapsible%closed]
```

```
//Start collapsible Definitions block
```

```
=====
```

```
[#href]
```

```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#unsupported_reason]
```

```
[.api-collapsible-fifth-title]
```

```
unsupported_reason
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|If volume activity tracking is not supported on the volume, this field provides an appropriate error code.
```

```
|message
```

```
|string
```

```
a|If volume activity tracking is not supported on the volume, this field provides an error message detailing why this is the case.
```

```
|===
```

```
[#activity_tracking]
```

```
[.api-collapsible-fifth-title]
```

```
activity_tracking
```

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|state
|string
a|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
a|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
|link:#unsupported_reason[unsupported_reason]
a|

|===

[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]

```

```

a|

|name
|string
a|

|uuid
|string
a|

|===

[#initialization]
[.api-collapsible-fifth-title]
initialization

[cols=3*,options=header]
|===
|Name
|Type
|Description

|state
|string
a|State of the analytics file system scan.

|===

[#unsupported_reason]
[.api-collapsible-fifth-title]
unsupported_reason

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|If file system analytics is not supported on the volume, this field
provides the error code explaining why.

```

```
|message
|string
a|If file system analytics is not supported on the volume, this field
provides the error message explaining why.
```

```
|===
```

```
[#analytics]
[.api-collapsible-fifth-title]
analytics
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|initialization
|link:#initialization[initialization]
a|
```

```
|scan_progress
|integer
a|Percentage of files in the volume that the file system analytics
initialization scan has processed. Only returned when the state is
`initializing`.
```

```
|state
|string
a|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
```

```
a|This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field.
```

```
|unsupported_reason
```

```
|link:#unsupported_reason[unsupported_reason]
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|suspects
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#anti_ransomware_attack_report]
```

```
[.api-collapsible-fifth-title]
```

```
anti_ransomware_attack_report
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```



```
|time
|string
a|Timestamp at which ransomware attack is observed.
```

```
|===
```

```
[#space]
[.api-collapsible-fifth-title]
space
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|snapshot_count
|integer
a|Total number of Anti-ransomware backup Snapshot copies.
```

```
|used
|integer
a|Total space in bytes used by the Anti-ransomware feature.
```

```
|used_by_logs
|integer
a|Space in bytes used by the Anti-ransomware analytics logs.
```

```
|used_by_snapshots
|integer
a|Space in bytes used by the Anti-ransomware backup Snapshot copies.
```

```
|===
```

```
[#suspect_files]
[.api-collapsible-fifth-title]
suspect_files
```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|count
|integer
a|Total number of `suspect_files.format` files observed by the Anti-
ransomware analytics engine on the volume.

|entropy
|string
a|Indicates the entropy level of this file type.

|format
|string
a|File formats observed by the Anti-ransomware analytics engine on the
volume.

```

```
|===
```

```

[#anti_ransomware]
[.api-collapsible-fifth-title]
anti_ransomware

```

Anti-ransomware related information of the volume.

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

```

```
|attack_probability
```

```
|string
```

```
a|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[link:#anti_ransomware_attack_report[anti_ransomware_attack_report]]
a|

|dry_run_start_time
|string
a|Time when Anti-ransomware monitoring `state` is set to dry-run value for
starting evaluation mode.

|space
|link:#space[space]
a|

|state
|string
a|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
a|Indicates whether or not to set the surge values as historical values.

|suspect_files
|array[link:#suspect_files[suspect_files]]
a|

|===

[#application]

```

```
[.api-collapsible-fifth-title]
```

```
application
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|Name of the application to which the volume belongs. Available only when the volume is part of an application.
```

```
|uuid
```

```
|string
```

```
a|UUID of the application to which the volume belongs. Available only when the volume is part of an application.
```

```
|===
```

```
[#asynchronous_directory_delete]
```

```
[.api-collapsible-fifth-title]
```

```
asynchronous_directory_delete
```

Configuration for asynchronous directory delete from the client. This is only supported on Flexible volumes and FlexGroup volumes.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

```
a|Specifies whether asynchronous directory delete from the client is enabled on the volume.
```

```
|trash_bin
```

```
|string
```

```
a|Name of the trash bin directory. If no "trash_bin" property is specified
```

when enabling, the default trash bin name, `"._ontaptrashbin"`, is used.

|===

```
[#autosize]
[.api-collapsible-fifth-title]
autosize
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|grow_threshold
```

```
|integer
```

a|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
```

a|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
```

a|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
```

a|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.

```
|shrink_threshold
|integer
a|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]
[.api-collapsible-fifth-title]
snapshot_reference
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#parent_svm]
[.api-collapsible-fifth-title]
parent_svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```

|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

```

```
|===
```

```

[#parent_volume]
[.api-collapsible-fifth-title]
parent_volume

```

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

```

```

|_links
|link:#_links[_links]
a|

```

```

|name
|string
a|The name of the volume.

```

```

|uuid
|string
a|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]
[.api-collapsible-fifth-title]
clone

[cols=3*,options=header]
|===
|Name
|Type
|Description

|inherited_physical_used
|integer
a|Inherited physical used from the clone's base snapshot.

|inherited_savings
|integer
a|Inherited savings from the clone's base snapshot.

|is_flexclone
|boolean
a|Specifies if this volume is a normal FlexVol or FlexClone. This field
needs to be set when creating a FlexClone. Valid in POST.

|parent_snapshot
|link:#snapshot_reference[snapshot_reference]
a|

|parent_svm
|link:#parent_svm[parent_svm]
a|

|parent_volume
|link:#parent_volume[parent_volume]
a|

|split_complete_percent
|integer
a|Percentage of FlexClone blocks split from its parent volume.

```


|split_estimate
|integer
a|Space required by the containing-aggregate to split the FlexClone volume.

|split_initiated
|boolean
a|This field is set when split is executed on any FlexClone, that is when the FlexClone volume is split from its parent FlexVol. This field needs to be set for splitting a FlexClone form FlexVol. Valid in PATCH.

|===

[#consistency_group]
[.api-collapsible-fifth-title]
consistency_group

Consistency group the volume is part of.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|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

a|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]
[.api-collapsible-fifth-title]
aggregates

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the aggregate hosting the FlexGroup Constituent.

|uuid
|string
a|Unique identifier for the aggregate.

|===

[#destination_aggregate]
[.api-collapsible-fifth-title]
destination_aggregate

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

```

```
|uuid
|string
a|
```

```
|===
```

```
[#movement]
[.api-collapsible-fifth-title]
movement
```

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|cutover_window
|integer
```

a|Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.

```
|destination_aggregate
|link:#destination_aggregate[destination_aggregate]
a|Aggregate
```

```
|percent_complete
|integer
a|Completion percentage
```

```
|state
|string
a|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

a|Tiering policy for FabricPool

|===

[#logical_space]

[.api-collapsible-fifth-title]

logical_space

[cols=3*,options=header]

|===

|Name

|Type

|Description

|available

|integer

a|The amount of space available in this volume with storage efficiency space considered used, in bytes.

|enforcement

|boolean

a|Specifies whether space accounting for operations on the volume is done along with storage efficiency.

|reporting

|boolean

a|Specifies whether space reporting on the volume is done along with storage efficiency.

|used_by_afs

|integer

a|The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes.

|===

```

[#snapshot]
[.api-collapsible-fifth-title]
snapshot

[cols=3*,options=header]
|===
|Name
|Type
|Description

|autodelete_enabled
|boolean
a|Specifies whether Snapshot copy autodelete is currently enabled on this
volume.

|reserve_percent
|integer
a|The space that has been set aside as a reserve for Snapshot copy usage,
in percent.

|used
|integer
a|The total space used by Snapshot copies in the volume, in bytes.

|===

[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]
|===
|Name
|Type
|Description

|afs_total
|integer
a|Total size of AFS, excluding snap-reserve, in bytes.

|available

```

```
|integer
a|The available space, in bytes.

|available_percent
|integer
a|The space available, as a percent.

|block_storage_inactive_user_data
|integer
a|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
a|Space used by capacity tier for this volume in the FabricPool aggregate,
in bytes.

|footprint
|integer
a|Data used for this volume in the aggregate, in bytes.

|large_size_enabled
|boolean
a|Specifies whether the support for large volumes and large files is
enabled on the volume.

|local_tier_footprint
|integer
a|Space used by the local tier for this volume in the aggregate, in bytes.

|logical_space
|link:#logical_space[logical_space]
a|

|metadata
|integer
a|Space used by the volume metadata in the aggregate, in bytes.
```

|over_provisioned
|integer
a|The amount of space not available for this volume in the aggregate, in bytes.

|performance_tier_footprint
|integer
a|Space used by the performance tier for this volume in the FabricPool aggregate, in bytes.

|size
|integer
a|Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes.

|snapshot
|link:#snapshot[snapshot]
a|

|total_footprint
|integer
a|Data and metadata used for this volume in the aggregate, in bytes.

|used
|integer
a|The virtual space used (includes volume reserves) before storage efficiency, in bytes.

|used_by_afs
|integer
a|The space used by Active Filesystem, in bytes.

|used_percent
|integer
a|The virtual space used (includes volume reserves) before storage efficiency, as a percent.

|===

```
[#constituents]
[.api-collapsible-fifth-title]
constituents
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|aggregates
```

```
|link:#aggregates[aggregates]
```

```
a|
```

```
|movement
```

```
|link:#movement[movement]
```

```
a|Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
```

```
|name
```

```
|string
```

```
a|FlexGroup Constituents name.
```

```
|space
```

```
|link:#space[space]
```

```
a|
```

```
|===
```

```
[#idcs_scanner]
```

```
[.api-collapsible-fifth-title]
```

```
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.

```
[cols=3*,options=header]
```



```
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies the administrative state of the inactive data compression
scanner.

|inactive_days
|integer
a|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
a|Specifies the mode of inactive data compression scanner. Valid for PATCH
and GET.

|operation_state
|string
a|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
a|Status of last inactive data compression scan on the volume.

|threshold_inactive_time
|string
a|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.+++</num>+++

|===
```

```

[#policy]
[.api-collapsible-fifth-title]
policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Specifies the name of the efficiency policy.

|===

[#scanner]
[.api-collapsible-fifth-title]
scanner

[cols=3*,options=header]
|===
|Name
|Type
|Description

|compression
|boolean
a|Start compression if scanning old data. Valid for PATCH and GET. This
option is not supported for FSX/CVO platforms.

|dedupe
|boolean
a|Start deduplication if scanning old data. Valid for PATCH and GET.

|scan_old_data
|boolean
a|Indicates whether or not to scan old data. Valid for PATCH and GET.

|state
|string
a|State of the volume efficiency scanner. Valid for PATCH and GET. Valid

```

options for PATCH are "idle" and "active".

|===

[#space_savings]

[.api-collapsible-fifth-title]

space_savings

[cols=3*,options=header]

|===

|Name

|Type

|Description

|compression

|integer

a|Total disk space that is saved by compressing blocks on the referenced file system, in bytes.

|compression_percent

|integer

a|Percentage of total disk space that is saved by compressing blocks on the referenced file system.

|dedupe

|integer

a|Total disk space that is saved by deduplication and file cloning, in bytes.

|dedupe_percent

|integer

a|Percentage of total disk space that is saved by deduplication and file cloning.

|dedupe_sharing

|integer

a|Total disk space that is shared due to deduplication and file cloning.

|total

|integer

a|Total disk space saved in the volume due to deduplication, compression and file cloning, in bytes.

|total_percent

|integer

a|Percentage of total disk space saved in the volume due to deduplication, compression and file cloning.

|===

[#efficiency]

[.api-collapsible-fifth-title]

efficiency

[cols=3*,options=header]

|===

|Name

|Type

|Description

|application_io_size

|string

a|Block size to use by compression.

|auto_state

|string

a|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

a|The system can be enabled/disabled compaction.

inline ‐ Data will be compacted first and written to the volume.

none ‐ None

mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.

|compression

|string

a|The system can be enabled/disabled compression.

inline ‐ Data will be compressed first and written to the volume.

background ‐ Data will be written to the volume and compressed later.

both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run.

none ‐ None

mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.

NOTE: that On volumes with container compression enabled, background compression refers to inactive data compression scan enabled on the volume.

|compression_type

|string

a|Compression type to use by compression. Valid for PATCH and GET.

|cross_volume_dedupe

|string

a|The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled.

inline ‐ Data will be cross volume deduped first and written to the volume.

background ‐ Data will be written to the volume and cross volume deduped later.

both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run.

none ‐ None

mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.

|dedupe

|string

a|The system can be enabled/disabled dedupe.

inline ‐ Data will be deduped first and written to the volume.

background ‐ Data will be written to the volume and deduped later.

both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run.

none ‐ None

mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.

|has_savings

|boolean

a|When true, indicates that the volume contains shared(deduplication, file clones) or compressed data.

|idcs_scanner

|link:#idcs_scanner[idcs_scanner]

a|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

a|Last sis operation begin timestamp.

|last_op_end

|string

a|Last sis operation end timestamp.

|last_op_err

|string

a|Last sis operation error text.

|last_op_size

|integer

a|Last sis operation size.

|last_op_state

|string

a|Last sis operation state.

|logging_enabled

|boolean

a|When true, indicates that space savings for any newly-written data are

being logged.

```
|op_state
|string
a|Sis status of the volume.
```

```
|policy
|link:#policy[policy]
a|
```

```
|progress
|string
a|Sis progress of the volume.
```

```
|scanner
|link:#scanner[scanner]
a|
```

```
|schedule
|string
a|Schedule associated with volume.
```

```
|space_savings
|link:#space_savings[space_savings]
a|
```

```
|state
|string
a|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 &dash; All storage efficiency features are disabled.
mixed &dash; Read-only field for FlexGroup volumes, storage efficiency is
enabled on certain constituents and disabled on others.
On FSx for ONTAP and Cloud Volumes ONTAP &dash;
    &emsp; enabled &dash; All supported storage efficiency features for the
volume are enabled.
    &emsp; custom &dash; 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 &dash;
    &emsp; enabled &dash; 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
```

```
a|Storage efficiency mode used by volume. This parameter is supported only on AFF platform.
```

```
|type
```

```
|string
```

```
a|Sis Type of the volume.
```

```
|volume_path
```

```
|string
```

```
a|Absolute volume path of the volume.
```

```
|===
```

```
[#status]
```

```
[.api-collapsible-fifth-title]
```

```
status
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Encryption progress message code.
```

```
|message
```

```
|string
```

```
a|Encryption progress message.
```

```
|===
```



```
[#encryption]
[.api-collapsible-fifth-title]
encryption
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

```
a|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
```

```
a|Encryption key creation time of the volume.
```

```
|key_id
```

```
|string
```

```
a|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
```

```
a|Specifies an additional key manager attribute that is an identifier-value pair, separated by '='. For example, CRN=unique-value. This parameter is required when using the POST method and an IBM Key Lore key manager is configured on the SVM.
```

```
|rekey
```

```
|boolean
```

```
a|If set to 'true', re-encrypts the volume with a new key. Valid in PATCH.
```

```

|state
|string
a|Volume encryption state.
encrypted &dash; The volume is completely encrypted.
encrypting &dash; Encryption operation is in progress.
partial &dash; Some constituents are encrypted and some are not.
Applicable only for FlexGroup volume.
rekeying. Encryption of volume with a new key is in progress.
unencrypted &dash; The volume is a plain-text one.

|status
|link:#status[status]
a|

|type
|string
a|Volume encryption type.
none &dash; The volume is a plain-text one.
volume &dash; The volume is encrypted with NVE (NetApp Volume Encryption).
aggregate &dash; The volume is encrypted with NAE (NetApp Aggregate
Encryption).

|===

[#error_state]
[.api-collapsible-fifth-title]
error_state

[cols=3*,options=header]
|===
|Name
|Type
|Description

|has_bad_blocks
|boolean
a|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
a|Indicates whether the file system has any inconsistencies.

```

```
true &dash; File system is inconsistent.  
false &dash; File system in not inconsistent.
```

```
|===
```

```
[#files]  
[.api-collapsible-fifth-title]  
files
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|maximum
```

```
|integer
```

a|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
```

a|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]  
[.api-collapsible-fifth-title]  
flash_pool
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

|cache_eligibility

|string

a|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

a|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

a|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]

[.api-collapsible-fifth-title]

flexgroup

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|Name of the FlexGroup volume that the constituent is part of.

|uuid

|string

a|Unique identifier for the FlexGroup volume that the constituent is part

of.

|===

```
[#guarantee]
[.api-collapsible-fifth-title]
guarantee
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

|honored

|boolean

a|Is the space guarantee of this volume honored in the aggregate?

|type

|string

a|The type of space guarantee of this volume in the aggregate.

|===

```
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#cloud]
[.api-collapsible-fifth-title]
cloud
```

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|duration
|string
```

a|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
|link:#iops[iops]
```

a|The rate of I/O operations observed at the storage object.

```
|latency
|link:#latency[latency]
```

a|The round trip latency in microseconds observed at the storage object.

```
|status
|string
```

a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#flexcache]
[.api-collapsible-fifth-title]
flexcache
```

Performance number for FlexCache used to measure cache effectiveness.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|bandwidth_savings
|integer
a|Bandwidth savings denoting the amount of data served locally by the
cache, in bytes.
```

```
|cache_miss_percent
|integer
a|Cache miss percentage.
```

```
|duration
|string
a|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
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

Performance numbers, such as IOPS, latency and throughput.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|cloud
|link:#cloud[cloud]
a|Performance numbers (IOPS and latency) for cloud store. These numbers
are relevant only for volumes hosted on FabricPools.
```

```
|duration
|string
a|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
|link:#flexcache[flexcache]
a|Performance number for FlexCache used to measure cache effectiveness.
```

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|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
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#movement]
[.api-collapsible-fifth-title]
movement
```

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cutover_window
|integer
a|Time window in seconds for cutover. The allowed range is between 30 to
300 seconds.

|destination_aggregate
|link:#destination_aggregate[destination_aggregate]
a|Aggregate

|percent_complete
|integer
a|Completion percentage

|start_time
|string
a|Start time of volume move.

|state
|string
a|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
a|Tiering policy for FabricPool

|===

```

```
[#export_policy]
[.api-collapsible-fifth-title]
export_policy
```

Export Policy

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|id
```

```
|integer
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|===
```

```
[#junction_parent]
```

```
[.api-collapsible-fifth-title]
```

```
junction_parent
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|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
a|Unique identifier for the parent volume.

|===

[#nas]
[.api-collapsible-fifth-title]
nas

[cols=3*,options=header]

|===

|Name
|Type
|Description

|export_policy
|link:#export_policy[export_policy]
a|Export Policy

|gid
|integer
a|The UNIX group ID of the volume. Valid in POST or PATCH.

|junction_parent
|link:#junction_parent[junction_parent]
a|

|path
|string
a|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
a|Security style associated with the volume. Valid in POST or PATCH.
mixed &dash; Mixed-style security
ntfs &dash; NTFS/Windows-style security
unified &dash; Unified-style security, unified UNIX, NFS and CIFS
permissions
unix &dash; Unix-style security.

|uid
|integer
a|The UNIX user ID of the volume. Valid in POST or PATCH.

|unix_permissions
|integer
a|UNIX permissions to be viewed as an octal number. It consists of 4
digits derived by adding up bits 4 (read), 2 (write) and 1 (execute).
First digit selects the set user ID(4), set group ID (2) and sticky (1)
attributes. The second digit selects permission for the owner of the file;
the third selects permissions for other users in the same group; the
fourth for other users not in the group. Valid in POST or PATCH. For
security style "mixed" or "unix", the default setting is 0755 in octal
(493 in decimal) and for security style "ntfs", the default setting is
0000. In cases where only owner, group and other permissions are given (as
in 755, representing the second, third and fourth digit), first digit is
assumed to be zero.

|===

[#policy]
[.api-collapsible-fifth-title]
policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

|max_throughput_iops
|integer
a|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
a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

|min_throughput_iops
|integer
a|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
a|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
a|The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.

|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

|===

[#qos]
[.api-collapsible-fifth-title]
qos

QoS information


```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|policy
```

```
|link:#policy[policy]
```

```
a|
```

```
|===
```

```
[#quota]
```

```
[.api-collapsible-fifth-title]
```

```
quota
```

Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

a|This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVols or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVols or FlexGroup volumes when the quota state is "off".

```
|state
```

```
|string
```

a|Quota state of the volume

```
|===
```

```
[#last_error]
```

```
[.api-collapsible-fifth-title]
```

```
last_error
```

Error information for the last failed file move on the constituent.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|integer
```

```
a|Error code of the last file move error on the constiuent.
```

```
|destination
```

```
|integer
```

```
a|DSID of the destination constituent of the last file move error on the  
constiuent.
```

```
|file_id
```

```
|integer
```

```
a|File ID of the last file move error on the constiuent.
```

```
|time
```

```
|string
```

```
a|Time of the last file move error on the constiuent.
```

```
|===
```

```
[#movement]
```

```
[.api-collapsible-fifth-title]
```

```
movement
```

Properties on this constituent related to file movement.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|file_moves_started
|integer
a|Number of file moves started on this constituent.
```

```
|last_error
|link:#last_error[last_error]
a|Error information for the last failed file move on the constituent.
```

```
|most_recent_start_time
|string
a|Start time of the most recent file move on the constiuent.
```

```
|===
```

```
[#blocks_skipped]
[.api-collapsible-fifth-title]
blocks_skipped
```

Number of blocks skipped by the scanner on this constiuent due to various reasons.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|efficiency_blocks
|integer
a|Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.
```

```
|efficiency_percent
|integer
a|Number of blocks skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.
```

```
|fast_truncate
|integer
a|Number of blocks skipped by the scanner on this constituent because fast
```

truncate is currently running on files.

|footprint_invalid

|integer

a|Number of blocks skipped by the scanner on this constituent because of files with invalid space footprints.

|in_snapshot

|integer

a|Number of blocks skipped by the scanner on this constituent because of files in Snapshot copies.

|incompatible

|integer

a|Number of blocks skipped by the scanner on this constituent because of incompatible files.

|metadata

|integer

a|Number of blocks skipped by the scanner on this constituent because of metadata files.

|on_demand_destination

|integer

a|Number of blocks skipped by the scanner on this constituent because of on demand destination files.

|other

|integer

a|Number of blocks skipped by the scanner on this constituent for all other reasons.

|remote_cache

|integer

a|Number of blocks skipped by the scanner on this constituent because of remote caches.

|too_large

|integer

a|Number of blocks skipped by the scanner on this constituent because of files that are larger than `rebalancing.max_file_size`.

|too_small

|integer

a|Number of blocks skipped by the scanner on this constituent because of files that are smaller than `rebalancing.min_file_size`.

|write_fenced

|integer

a|Number of blocks skipped by the scanner on this constituent because of files fenced for write operations.

|===

[#files_skipped]

[.api-collapsible-fifth-title]

files_skipped

Number of files skipped by the scanner on this constituent due to various reasons.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|efficiency_blocks

|integer

a|Number of files skipped by the scanner on this constituent because storage efficiency lost, in blocks, would be too high.

|efficiency_percent

|integer

a|Number of files skipped by the scanner on this constituent because storage efficiency lost, in percent, would be too high.

|fast_truncate

|integer

a|Number of files skipped by the scanner on this constituent because fast truncate is currently running on the file.

|footprint_invalid

|integer

a|Number of files skipped by the scanner on this constituent because their space footprints are invalid.

|in_snapshot

|integer

a|Number of files skipped by the scanner on this constituent because they are trapped in Snapshot copies.

|incompatible

|integer

a|Number of files skipped by the scanner on this constituent because they are incompatible.

|metadata

|integer

a|Number of files skipped by the scanner on this constituent because they metadata files.

|on_demand_destination

|integer

a|Number of files skipped by the scanner on this constituent because they are on demand destinations.

|other

|integer

a|Number of files skipped by the scanner on this constituent for all other reasons.

|remote_cache

|integer

a|Number of files skipped by the scanner on this constituent because they are remote caches.

|too_large

```

|integer
a|Number of files skipped by the scanner on this constituent because they
are larger than rebalancing.max_file_size.

|too_small
|integer
a|Number of files skipped by the scanner on this constituent because they
are smaller than rebalancing.min_file_size.

|write_fenced
|integer
a|Number of files skipped by the scanner on this constituent because they
are fenced for write operations.

|===

[#scanner]
[.api-collapsible-fifth-title]
scanner

Properties related to determining which files to move and where to move
them to.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|blocks_scanned
|integer
a|Number of blocks scanned on this constituent.

|blocks_skipped
|link:#blocks_skipped[blocks_skipped]
a|Number of blocks skipped by the scanner on this constiuent due to
various reasons.

|files_scanned
|integer

```

```
a|Number of files scanned on this constituent.
```

```
|files_skipped
```

```
|link:#files_skipped[files_skipped]
```

```
a|Number of files skipped by the scanner on this constituent due to various reasons.
```

```
|===
```

```
[#engine]
```

```
[.api-collapsible-fifth-title]
```

```
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|movement
```

```
|link:#movement[movement]
```

```
a|Properties on this constituent related to file movement.
```

```
|scanner
```

```
|link:#scanner[scanner]
```

```
a|Properties related to determining which files to move and where to move them to.
```

```
|===
```

```
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
```

```
error_arguments
```



```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
[#rebalancing]  
[.api-collapsible-fifth-title]  
rebalancing
```

Configuration and runtime properties involving non-disruptive volume capacity rebalancing for a FlexGroup volume.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|data_moved
```

```
|integer
```

a|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
```

```
|link:#engine[engine]
```

a|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
```

a|Specifies whether or not to exclude files that are stuck in Snapshot copies 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

a|Represents the percentage the volume is out of balance.

|imbalance_size

|integer

a|Represents how much the volume is out of balance, in bytes.

|max_constituent_imbalance_percent

|integer

a|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

a|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

a|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>++++

|max_threshold

|integer

a|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

a|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.

|min_threshold

|integer

a|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[link:#error[error]]

a|Capacity rebalancing notice messages.

|runtime

|string

a|Duration the capacity rebalancing operation has been running.

|start_time

|string

a|Time when the current capacity rebalancing operation started, or when a future scheduled rebalancing operation begins.

|state

|string

a|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.

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.

idle ‐ capacity rebalancing is running on the constituent, however, no active scanning or file movement is currently occurring.

scanning ‐ the constituent's file system is being scanned to find files to move and determine free space.

rebalancing_source ‐ a file is being moved off of the constituent.

rebalancing_dest ‐ a file is being moved to the constituent.

not_running ‐ capacity rebalancing is not running on the constituent.

unknown ‐ the system was unable to determine the rebalancing state for the constituent.

|stop_time

|string

a|Time when the capacity rebalancing operation stopped.

|target_used

|integer

a|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

a|Represents the used size of each constituent, as determined by the rebalancing engine. Calculated by subtracting the size used by Snapshot copies, the size of files pending deletion and the size of filesystem metadata from the volume used size.

|===

[#retention]

[.api-collapsible-fifth-title]

retention

[cols=3*,options=header]

|===

|Name

|Type

|Description

|default

|string

a|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>+++

|maximum

|string

a|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>+++

|minimum

|string

a|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]
```

```
[.api-collapsible-fifth-title]
```

```
snaplock
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|append_mode_enabled
```

```
|boolean
```

a|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.

```
|autocommit_period
```

```
|string
```

a|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
```

a|This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.


```
|expiry_time
|string
a|Expiry time of the volume.

|is_audit_log
|boolean
a|Indicates if this volume has been configured as SnapLock audit log
volume for the SVM .

|litigation_count
|integer
a|Litigation count indicates the number of active legal-holds on the
volume.

|privileged_delete
|string
a|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
|link:#retention[retention]
a|

|type
|string
a|The SnapLock type of the volume.
compliance &dash; A SnapLock Compliance(SLC) volume provides the highest
level of WORM protection and an administrator cannot destroy a SLC volume
if it contains unexpired WORM files.
enterprise &dash; An administrator can delete a SnapLock Enterprise(SLE)
volume.
non_snaplock &dash; Indicates the volume is non-snaplock.

|unspecified_retention_file_count
|integer
a|Indicates the number of files with an unspecified retention time in the
volume.
```

```
|===
```

```
[#destinations]  
[.api-collapsible-fifth-title]  
destinations
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|is_cloud
```

```
|boolean
```

```
a|Specifies whether a volume is a SnapMirror source volume, using  
SnapMirror to protect its data to a cloud destination.
```

```
|is_ontap
```

```
|boolean
```

```
a|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]  
[.api-collapsible-fifth-title]  
snapmirror
```

```
Specifies attributes for SnapMirror protection.
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```

|destinations
|link:#destinations[destinations]
a|

|is_protected
|boolean
a|Specifies whether a volume is a SnapMirror source volume, using
SnapMirror to protect its data.

|===

[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy

This is a reference to the Snapshot copy policy.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#logical_space]
[.api-collapsible-fifth-title]
logical_space

[cols=3*,options=header]
|===

```

Name	Type	Description
available	integer	a The amount of space available in this volume with storage efficiency space considered used, in bytes.
enforcement	boolean	a Specifies whether space accounting for operations on the volume is done along with storage efficiency.
reporting	boolean	a Specifies whether space reporting on the volume is done along with storage efficiency.
used	integer	a SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), in bytes.
used_by_afs	integer	a The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes.
used_by_snapshots	integer	a Size that is logically used across all Snapshot copies in the volume, in bytes.
used_percent	integer	a SUM of (physical-used, shared_refs, compression_saved_in_plane0, vbn_zero, future_blk_cnt), as a percentage.
===		

```

[#autodelete]
[.api-collapsible-fifth-title]
autodelete

[cols=3*,options=header]
|===
|Name
|Type
|Description

|commitment
|string
a|By default, Snapshot copy autodelete does not delete Snapshot copies
locked by Snapmirror, clones of a volume, a LUN, an NVMe namespace, or a
file. Deletion of Snapshot copies locked by these applications is
specified using this option. The default value is try.

|defer_delete
|string
a|Allows the user to inform Snapshot copy autodelete to defer the deletion
of a specified Snapshot copy until the end. The default value is
user_created.

|delete_order
|string
a|Specifies the order in which Snapshot copy autodelete occurs. Ordering
is done using the date and time the Snapshot copy is created. The default
value is oldest_first.

|enabled
|boolean
a|Specifies whether Snapshot copy autodelete is currently enabled on this
volume.

|prefix
|string
a|Specifies the prefix of the Snapshot copy which if matched, is deleted
last. Used with autodelete_defer_delete when used with a prefix value.

|target_free_space

```

```
|integer
a|Snapshot copies 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
a|Specifies when the system should trigger an autodelete of Snapshot
copies. When set to _volume_, autodelete is triggered based on volume
fullness. When set to _snap_reserve_, autodelete is triggered based on
Snapshot copy reserve fullness. The default value is _volume_.
```

```
|===
```

```
[#snapshot]
[.api-collapsible-fifth-title]
snapshot
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|autodelete
|link:#autodelete[autodelete]
a|
```

```
|autodelete_enabled
|boolean
a|Specifies whether Snapshot copy 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
a|Specifies when the system should trigger an autodelete of Snapshot
copies. When set to _volume_, autodelete is triggered based on volume
fullness. When set to _snap_reserve_, autodelete is triggered based on
Snapshot copy 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
a|Size available for Snapshot copies within the Snapshot copy reserve, in bytes.

|reserve_percent
|integer
a|The space that has been set aside as a reserve for Snapshot copy usage, in percent.

|reserve_size
|integer
a|Size in the volume that has been set aside as a reserve for Snapshot copy usage, in bytes.

|space_used_percent
|integer
a|Percentage of snapshot reserve size that has been used.

|used
|integer
a|The total space used by Snapshot copies in the volume, in bytes.

|===

[#space]
[.api-collapsible-fifth-title]
space

[cols=3*,options=header]

|===

|Name
|Type
|Description

|afs_total
|integer
a|Total size of AFS, excluding snap-reserve, in bytes.

|auto_adaptive_compression_footprint_data_reduction

|integer

a|Savings achieved due to Auto Adaptive Compression, in bytes.

|available

|integer

a|The available space, in bytes.

|available_percent

|integer

a|The space available, as a percent.

|block_storage_inactive_user_data

|integer

a|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

a|Percentage of size that is physically used in the performance tier of the volume.

|capacity_tier_footprint

|integer

a|Space used by capacity tier for this volume in the FabricPool aggregate, in bytes.

|capacity_tier_footprint_data_reduction

|integer

a|Savings achieved in the space used by the capacity tier for this volume in the FabricPool aggregate, in bytes.

|cross_volume_dedupe_metafiles_footprint

|integer

a|Cross volume deduplication metadata footprint, in bytes.

|cross_volume_dedupe_metafiles_temporary_footprint

|integer
a|Cross volume temporary deduplication metadata footprint, in bytes.

|dedupe_metafiles_footprint
|integer
a|Deduplication metadata footprint, in bytes.

|dedupe_metafiles_temporary_footprint
|integer
a|Temporary deduplication metadata footprint, in bytes.

|delayed_free_footprint
|integer
a|Delayed free blocks footprint, in bytes.

|effective_total_footprint
|integer
a|Volume footprint after efficiency savings, in bytes.

|expected_available
|integer
a|Size that should be available for the volume, irrespective of available size in the aggregate, in bytes.

|file_operation_metadata
|integer
a|File operation metadata footprint, in bytes.

|filesystem_size
|integer
a|Total usable size of the volume, in bytes.

|filesystem_size_fixed
|boolean
a|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
a|Data used for this volume in the aggregate, in bytes.

|fractional_reserve
|integer
a|Used to change the amount of space reserved for overwrites of reserved
objects in a volume.

|full_threshold_percent
|integer
a|Volume full threshold percentage at which EMS warnings can be sent.

|is_used_stale
|boolean
a|Specifies if the virtual space used is stale.

|large_size_enabled
|boolean
a|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
a|Space used by the local tier for this volume in the aggregate, in bytes.

|logical_space
|link:#logical_space[logical_space]
a|

|metadata
|integer
a|Space used by the volume metadata in the aggregate, in bytes.

|nearly_full_threshold_percent
|integer
a|Volume nearly full threshold percentage at which EMS warnings can be
sent.
```

|over_provisioned
|integer
a|The amount of space not available for this volume in the aggregate, in bytes.

|overwrite_reserve
|integer
a|Reserved space for overwrites, in bytes.

|overwrite_reserve_used
|integer
a|Overwrite logical reserve space used, in bytes.

|percent_used
|integer
a|Percentage of the volume size that is used.

|performance_tier_footprint
|integer
a|Space used by the performance tier for this volume in the FabricPool aggregate, in bytes.

|physical_used
|integer
a|Size that is physically used in the volume, in bytes.

|physical_used_percent
|integer
a|Size that is physically used in the volume, as a percentage.

|size
|integer
a|Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes.

|size_available_for_snapshots
|integer
a|Available space for Snapshot copies from snap-reserve, in bytes.

|snapmirror_destination_footprint
|integer
a|SnapMirror destination footprint, in bytes.

|snapshot
|link:#snapshot[snapshot]
a|

|snapshot_reserve_unusable
|integer
a|Snapshot reserve that is not available for Snapshot copy creation, in bytes.

|snapshot_spill
|integer
a|Space used by the snapshot copies beyond the snap-reserve, in bytes.

|total_footprint
|integer
a|Data and metadata used for this volume in the aggregate, in bytes.

|used
|integer
a|The virtual space used (includes volume reserves) before storage efficiency, in bytes.

|used_by_afs
|integer
a|The space used by Active Filesystem, in bytes.

|user_data
|integer
a|User data, in bytes.

|volume_guarantee_footprint
|integer
a|Space reserved for future writes in the volume, in bytes.

```
|===
```

```
[#access]  
[.api-collapsible-fifth-title]  
access
```

Raw count and latency data for access operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
```

a|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]  
[.api-collapsible-fifth-title]  
audit
```

Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
dir
```

```
Raw count and latency data for directory-create operations.
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]  
[.api-collapsible-fifth-title]  
file
```

Raw count and latency data for file-create operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
```

a|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]  
[.api-collapsible-fifth-title]  
other
```

Raw count and latency data for create operations on objects other than files, directories and symlinks.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
symlink
```

```
Raw count and latency data for symlink-create operations.
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]  
[.api-collapsible-fifth-title]  
create
```

Raw count and latency data for create operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|dir
```

```
|link:#dir[dir]
```

a|Raw count and latency data for directory-create operations.

```
|file
```

```
|link:#file[file]
```

a|Raw count and latency data for file-create operations.

```
|other
```

```
|link:#other[other]
```

a|Raw count and latency data for create operations on objects other than files, directories and symlinks.

```
|symlink
```

```
|link:#symlink[symlink]
```

a|Raw count and latency data for symlink-create operations.

```
|===
```

```
[#getattr]  
[.api-collapsible-fifth-title]  
getattr
```

Raw count and latency data for getattr operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
link
```

Raw count and latency data for link operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
[.api-collapsible-fifth-title]
lock
```

Raw count and latency data for lock operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|count
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
|integer
```

a|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]
[.api-collapsible-fifth-title]
lookup
```

Raw count and latency data for lookup operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
open
```

```
Raw count and latency data for open operations.
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
[.api-collapsible-fifth-title]
read
```

Raw count and latency data for read operations, including histograms categorizing operations by size and latency.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|count
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
|integer
```

a|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]
a|
```

```
|volume_protocol_latency_histogram_labels
|array[string]
```

a|Labels for the latency histogram, ranging from <2us to >20s.

```
|volume_protocol_size_histogram_counts
|array[integer]
a|
```

```
|volume_protocol_size_histogram_labels
|array[string]
a|Labels for the size histogram, ranging from <4KB to >1024KB.
```

```
|===
```

```
[#readdir]
[.api-collapsible-fifth-title]
readdir
```

Raw count and latency data for readdir operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|count
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
|integer
```

```
a|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]
[.api-collapsible-fifth-title]
readlink
```

Raw count and latency data for readlink operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
```

```
a|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]
```

```
[.api-collapsible-fifth-title]
```

```
rename
```

Raw count and latency data for rename operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|Number of operations of the given type performed on this volume.
```

```
|total_time
```

```
|integer
a|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]
[.api-collapsible-fifth-title]
setattr
```

Raw count and latency data for setattr operations.

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|count
|integer
a|Number of operations of the given type performed on this volume.
```

```
|total_time
|integer
a|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]
[.api-collapsible-fifth-title]
unlink
```


Raw count and latency data for unlink operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
```

a|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]
```

```
[.api-collapsible-fifth-title]
```

```
watch
```

Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

a|Number of operations of the given type performed on this volume.

```
|total_time
```

```
|integer
a|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]
[.api-collapsible-fifth-title]
write
```

Raw count and latency data for write operations, including histograms categorizing operations by size and latency.

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|count
|integer
a|Number of operations of the given type performed on this volume.
```

```
|total_time
|integer
a|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]
a|
```

```
|volume_protocol_latency_histogram_labels
|array[string]
a|Labels for the latency histogram, ranging from <2us to >20s.
```

```
|volume_protocol_size_histogram_counts
|array[integer]
a|
```

```
|volume_protocol_size_histogram_labels
|array[string]
a|Labels for the size histogram, ranging from <4KB to >1024KB.
```

```
|===
```

```
[#cifs_ops_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|access
```

```
|link:#access[access]
```

```
a|Raw count and latency data for access operations.
```

```
|audit
```

```
|link:#audit[audit]
```

```
a|Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.
```

```
|create
```

```
|link:#create[create]
```

```
a|Raw count and latency data for create operations.
```

```
|getattr
|link:#getattr[getattr]
a|Raw count and latency data for getattr operations.
```

```
|link
|link:#link[link]
a|Raw count and latency data for link operations.
```

```
|lock
|link:#lock[lock]
a|Raw count and latency data for lock operations.
```

```
|lookup
|link:#lookup[lookup]
a|Raw count and latency data for lookup operations.
```

```
|open
|link:#open[open]
a|Raw count and latency data for open operations.
```

```
|read
|link:#read[read]
a|Raw count and latency data for read operations, including histograms
categorizing operations by size and latency.
```

```
|readdir
|link:#readdir[readdir]
a|Raw count and latency data for readdir operations.
```

```
|readlink
|link:#readlink[readlink]
a|Raw count and latency data for readlink operations.
```

```
|rename
|link:#rename[rename]
a|Raw count and latency data for rename operations.
```

```
|setattr
|link:#setattr[setattr]
a|Raw count and latency data for setattr operations.
```

```
|unlink
|link:#unlink[unlink]
a|Raw count and latency data for unlink operations.
```

```
|watch
|link:#watch[watch]
a|Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.
```

```
|write
|link:#write[write]
a|Raw count and latency data for write operations, including histograms categorizing operations by size and latency.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#cloud]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

```
a|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
```

```
|link:#latency_raw[latency_raw]
```

```
a|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
```

```
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#flexcache_raw]
[.api-collapsible-fifth-title]
flexcache_raw
```

Performance numbers for FlexCache used to measure cache effectiveness.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|cache_miss_blocks
|integer
a|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
a|Total blocks requested by the client.
```

```
|status
|string
a|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  
a|The timestamp of the performance data.
```

```
|===
```

```
[#nfs_ops_raw]  
[.api-collapsible-fifth-title]  
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|access
```

```
|link:#access[access]
```

```
a|Raw count and latency data for access operations.
```

```
|audit
```

```
|link:#audit[audit]
```

```
a|Raw count and latency data for audit operations. These statistics are only applicable for CIFS protocol operations.
```

```
|create
```

```
|link:#create[create]
```

```
a|Raw count and latency data for create operations.
```

```
|getattr
|link:#getattr[getattr]
a|Raw count and latency data for getattr operations.
```

```
|link
|link:#link[link]
a|Raw count and latency data for link operations.
```

```
|lock
|link:#lock[lock]
a|Raw count and latency data for lock operations.
```

```
|lookup
|link:#lookup[lookup]
a|Raw count and latency data for lookup operations.
```

```
|open
|link:#open[open]
a|Raw count and latency data for open operations.
```

```
|read
|link:#read[read]
a|Raw count and latency data for read operations, including histograms
categorizing operations by size and latency.
```

```
|readdir
|link:#readdir[readdir]
a|Raw count and latency data for readdir operations.
```

```
|readlink
|link:#readlink[readlink]
a|Raw count and latency data for readlink operations.
```

```
|rename
|link:#rename[rename]
a|Raw count and latency data for rename operations.
```

```
|setattr
|link:#setattr[setattr]
a|Raw count and latency data for setattr operations.
```

```
|unlink
|link:#unlink[unlink]
a|Raw count and latency data for unlink operations.
```

```
|watch
|link:#watch[watch]
a|Raw count and latency data for watch operations. These statistics are only applicable for CIFS protocol operations.
```

```
|write
|link:#write[write]
a|Raw count and latency data for write operations, including histograms categorizing operations by size and latency.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
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.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|cifs_ops_raw
|link:#cifs_ops_raw[cifs_ops_raw]
a|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
|link:#cloud[cloud]
a|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
|link:#flexcache_raw[flexcache_raw]
a|Performance numbers for FlexCache used to measure cache effectiveness.

|iops_raw
|link:#iops_raw[iops_raw]
a|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
|link:#latency_raw[latency_raw]
a|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
|link:#nfs_ops_raw[nfs_ops_raw]
a|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
a|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
|link:#throughput_raw[throughput_raw]
a|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
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

SVM containing the volume. Required on POST.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#tiering]
[.api-collapsible-fifth-title]
tiering
```

```
[cols=3*,options=header]
```

|===

|Name

|Type

|Description

|min_cooling_days

|integer

a|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]

a|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".

|policy

|string

a|Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.

- all ‐ This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.
- auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store
- none ‐ Volume blocks will not be tiered to the cloud store.
- snapshot_only ‐ This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days.

```
|storage_class
|string
a|This parameter specifies the storage class that a FabricPool uses. This feature is only available on volumes in a FabricPools on FSx or Cloud Volumes ONTAP for AWS.
```

```
* Default value: 1
* enum: ["default", "S3_standard", "S3_standard_IA", "S3_glacier_IR"]
* Introduced in: 9.13
* x-ntap-modifyOnly: true
* x-nullable: true
```

```
|supported
|boolean
a|This parameter specifies whether or not FabricPools are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only FabricPool aggregates are used if this parameter is set to true and only non FabricPool aggregates are used if this parameter is set to false. Tiering support for a FlexGroup can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only does the volume need to support tiering by using FabricPools, the tiering "policy" must not be 'none'. A volume that uses FabricPools but has a tiering "policy" of 'none' supports tiering, but will not tier any data.
```

```
|===
```

```
[#volume]
[.api-collapsible-fifth-title]
volume
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|_tags
|array[string]
```


a|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

a|Indicates whether or not access time updates are enabled on the volume.

|activity_tracking

|link:#activity_tracking[activity_tracking]

a|

|aggregates

|array[link:#aggregates[aggregates]]

a|Aggregate hosting the volume. Required on POST.

|aggressive_readahead_mode

|string

a|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.

* Default value: 1

* enum: ["none", "file_prefetch"]

* Introduced in: 9.13

* x-nullable: true

|analytics

|link:#analytics[analytics]

a|

|anti_ransomware

|link:#anti_ransomware[anti_ransomware]

a|Anti-ransomware related information of the volume.

|anti_ransomware_state

|string

a|The Anti-ransomware state of the volume. If no "anti_ransomware_state" property is specified, the volume inherits the value from its parent SVM's "anti_ransomware_default_volume_state" property. If this value is "disabled", Anti-ransomware is disabled on the volume. If this value is

"enabled", Anti-ransomware is enabled on the volume and alerts are raised if any suspect is detected for those volumes. If this value is "dry_run", Anti-ransomware is enabled in the dry-run or learning mode on the volume. The "dry_run" state is same as the "enabled" state except that the analytics data is used here for learning. No alerts are raised for any detections or violations. If this value is "paused", Anti-ransomware is paused on the volume. Additionally, three more states are available, which are only valid for GET. If this value is "disable_in_progress", Anti-ransomware monitoring is being disabled and a cleanup operation is in effect. If this value is "enable_paused", Anti-ransomware is paused on the volume from its earlier enabled state. If this value is "dry_run_paused", Anti-ransomware monitoring is paused on the volume from its earlier dry_run state. For POST, the valid Anti-ransomware states are only "disabled", "enabled" and "dry_run", whereas for PATCH, "paused" is also valid along with the three valid states for POST.

```
|application
|link:#application[application]
a|
```

```
|asynchronous_directory_delete
|link:#asynchronous_directory_delete[asynchronous_directory_delete]
a|Configuration for asynchronous directory delete from the client. This is
only supported on Flexible volumes and FlexGroup volumes.
```

```
|autosize
|link:#autosize[autosize]
a|
```

```
|clone
|link:#clone[clone]
a|
```

```
|cloud_retrieval_policy
|string
a|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

a|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 on FSx or Cloud Volumes ONTAP.

* Introduced in: 9.13

* x-ntap-readModify: true

* x-nullable: true

|comment

|string

a|A comment for the volume. Valid in POST or PATCH.

|consistency_group

|link:#consistency_group[consistency_group]

a|Consistency group the volume is part of.

|constituents

|array[link:#constituents[constituents]]

a|FlexGroup Constituents. FlexGroup 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

a|Specifies the number of times to iterate over the aggregates listed with the "aggregates.name" or "aggregates.uuid" when creating or expanding a FlexGroup volume. If a volume is being created on a single aggregate, the system creates a flexible volume if the "constituents_per_aggregate" field

is not specified, or a FlexGroup volume if it is specified. If a volume is being created on multiple aggregates, the system always creates a FlexGroup volume. The root constituent of a FlexGroup volume is always placed on the first aggregate in the list, unless 'optimize_aggregates' is specified as 'true'.

|convert_unicode

|boolean

a|Specifies whether directory Unicode format conversion is enabled when directories are accessed by NFS clients.

|create_time

|string

a|Creation time of the volume. This field is generated when the volume is created.

|efficiency

|link:#efficiency[efficiency]

a|

|encryption

|link:#encryption[encryption]

a|

|error_state

|link:#error_state[error_state]

a|

|files

|link:#files[files]

a|

|flash_pool

|link:#flash_pool[flash_pool]

a|

|flexcache_endpoint_type

|string

a|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
|link:#flexgroup[flexgroup]
a|
```

```
|granular_data
|boolean
a|State of granular data on the volume. This setting is true by default when creating a new FlexGroup volume, but can be specified as false at the time of creation via a POST request. On FlexVol volumes, the setting is always false, as only FlexGroup volumes and FlexGroup constituents support this feature. Once enabled, this setting can only be disabled by restoring a Snapshot copy. Earlier versions of ONTAP (pre 9.11) 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 copy that was taken before the setting was enabled.
```

```
* Introduced in: 9.12
* x-nullable: true
```

```
|guarantee
|link:#guarantee[guarantee]
a|
```

```
|is_object_store
|boolean
a|Specifies whether the volume is provisioned for an object store server.
```

```
|is_svm_root
|boolean
a|Specifies whether the volume is a root volume of the SVM it belongs to.
```

```
|language
|string
a|Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
```

```
|max_dir_size
|integer
a|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
|link:#metric[metric]
a|Performance numbers, such as IOPS, latency and throughput.

|movement
|link:#movement[movement]
a|Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

|msid
|integer
a|The volume's Master Set ID.

|name
|string
a|Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.

|nas
|link:#nas[nas]
a|

|optimize_aggregates
|boolean
a|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
|link:#qos[qos]
a|QoS information

```
|queue_for_encryption
|boolean
a|Specifies whether the volume is queued for encryption.

|quota
|link:#quota[quota]
a|Quotas track the space or file usage of a user, group, or qtree in a
FlexVol or a FlexGroup volume.

|rebalancing
|link:#rebalancing[rebalancing]
a|Configuration and runtime properties involving non-disruptive volume
capacity rebalancing for a FlexGroup volume.

|scheduled_snapshot_naming_scheme
|string
a|Naming Scheme for automatic Snapshot copies:

* create_time - Automatic Snapshot copies are saved as per the start of
their current date and time.
* ordinal - Latest automatic snapshot copy is saved as
+++<scheduled_frequency>+++0 and subsequent copies will follow the
create_time naming convention.+++</scheduled_frequency>+++

|size
|integer
a|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
|link:#snaplock[snaplock]
a|

|snapmirror
|link:#snapmirror[snapmirror]
a|Specifies attributes for SnapMirror protection.

|snapshot_count
```

```
|integer
a|Number of Snapshot copies in the volume.

|snapshot_directory_access_enabled
|boolean
a|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
a|Specifies whether or not snapshot copy locking is enabled on the volume.

|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the Snapshot copy policy.

|space
|link:#space[space]
a|

|state
|string
a|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
|link:#statistics[statistics]
a|These are raw performance numbers, such as IOPS latency and throughput.
These numbers are aggregated across all nodes in the cluster and increase
with the uptime of the cluster.

|status
|array[string]
a|Describes the current status of a volume.
```


|style
|string
a|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 constituents.

|svm
|link:#svm[svm]
a|SVM containing the volume. Required on POST.

|tiering
|link:#tiering[tiering]
a|

|type
|string
a|Type of the volume.
rw ‐ read-write volume.
dp ‐ data-protection volume.
ls ‐ load-sharing `dp` volume. Valid in GET.

|use_mirrored_aggregates
|boolean
a|Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster

configuration.

|uuid

|string

a|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

|===

[#job_link]

[.api-collapsible-fifth-title]

job_link

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|uuid

|string

a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

|===

//end collapsible .Definitions block

====

:leveloffset: -1

= Manage files and directories

:leveloffset: +1

[[ID9cd28a28cf037dc5b452b783bab67136]]

= Storage volumes volume.uuid files path endpoint overview

== Overview

This API is used to read a file, write to a file, retrieve a list of files and directories, and retrieve or modify certain properties of files and directories. The path field is used to specify the path to the directory or file to be acted on. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided.

== File data

Read and write data from/to a named file. To read a file, the Accept request HTTP header must be specified as multipart/form-data, and a value for the `length` query property, which represents the number of bytes to be read, must be specified. The API will fail if the length of data being read/written exceeds 1 MB. This API should only be used on normal files or streams associated with files. The results for other file types, such as LUNs is undefined.

The following APIs are used to read or write data to a file:

– GET

```
/api/storage/volumes/{volume.uuid}/files/{path}?byte_offset=0&length=40  
-H "Accept: multipart/form-data"
```

– POST /api/storage/volumes/{volume.uuid}/files/{path} -H
"Content-Type: multipart/form-data" --form "file=the data to be written to
the new file"

– PATCH

```
/api/storage/volumes/{volume.uuid}/files/{path}?byte_offset=10 -H  
"Content-Type: multipart/form-data" --form "file=the new data to be  
written or overwritten to the existing file starting at byte_offset"
```

== Listing directories and files

A list of files and directories and their properties can be retrieved for a specified path.

The following APIs are used to view a list of files and directories:

– GET /api/storage/volumes/{volume.uuid}/files

– GET /api/storage/volumes/{volume.uuid}/files/{path}

– GET
/api/storage/volumes/{volume.uuid}/files/{path}?fields=*

== File information

The metadata and detailed information about a single directory or file can be retrieved by setting the `return_metadata` query property to `true`. The information returned includes `type`, `creation_time`, `modified_time`, `changed_time`, `accessed_time`, `unix_permissions`, `owner_id`, `group_id`, `size`, `hard_links_count`, `inode_number`, `is_empty`, `bytes_used`, `unique_bytes`, `inode_generation`, `is_vm_aligned`, `is_junction`, `links`, and `analytics` (if requested).

The following API is used to view the properties of a single file or directory:

– GET
/api/storage/volumes/{volume.uuid}/files/{path}?return_metadata=true

== File usage

Custom details about the usage of a file can be retrieved by specifying a value for the `byte_offset` and `length` query properties.

The following API is used to view the unique bytes, and bytes used, by a file based on the range defined by `byte_offset` and `length`:

– GET
/api/storage/volumes/{volume.uuid}/files/{path}?return_metadata=true&byte_offset={int}&length={int}

== Create a directory

The following API is used to create a directory:

– POST /api/storage/volumes/{volume.uuid}/files/{path} -d '{
"type" : "directory", "unix-permissions" : "644"}'

== Delete an entire directory

A directory can be deleted. The behavior of this call is equivalent to `rm -rf`.

The following API is used to delete an entire directory:

```
&ndash; DELETE  
/api/storage/volumes/{volume.uuid}/files/{path}?recurse=true
```

== Delete a file or an empty directory

The following API is used to delete a file or an empty directory:

```
&ndash; DELETE    /api/storage/volumes/{volume.uuid}/files/{path}
```

```
&ndash; DELETE  
/api/storage/volumes/{volume.uuid}/files/{path}?recurse=false
```

== File system analytics

File system analytics provide a quick method for obtaining information summarizing properties of all files within any directory tree of a volume. When file system analytics are enabled on a volume, ``analytics.+++`` fields may be requested, and will be populated in the response records corresponding to directories. The API does not support file system analytics for requests that are made beyond the boundary of the specified ``volume.uuid``.

The following APIs are used to obtain analytics information for a directory:

```
&ndash; GET  
/api/storage/volumes/{volume.uuid}/files/{path}?fields=analytics
```

```
&ndash; GET    /api/storage/volumes/{volume.uuid}/files/{path}?fields=**
```

== QoS

QoS policies and settings enforce Service Level Objectives (SLO) on a file. A pre-created QoS policy can be used by specifying the ``qos.name`` or ``qos.uuid`` properties.

The following APIs are used to assign a QoS policy to a file:

```
&ndash; PATCH    /api/storage/volumes/{volume.uuid}/files/{path} -d '{  
"qos_policy.name" : "policy" }'
```

```
&ndash; PATCH /api/storage/volumes/{volume.uuid}/files/{path} -d '{
"qos_policy.uuid" : "b89bc5dd-94a3-11e8-a7a3-0050568edf84" }'
```

== Symlinks

The following APIs are used to create a symlink and read the contents of a symlink:

```
&ndash; POST /api/storage/volumes/{volume.uuid}/files/{path} -d '{
"target" : "directory2/file1" }'
```

```
&ndash; GET
/api/storage/volumes/{volume.uuid}/files/{path}?return_metadata=true&fields=target
```

== Rename a file or a directory

The following API can be used to rename a file or a directory. Note that you need to provide the path relative to the root of the volume in the `path` body parameter.

```
&ndash; PATCH /api/storage/volumes/{volume.uuid}/files/{path} -d '{
"path" : "directory1/directory2" }'
```

```
&ndash; PATCH /api/storage/volumes/{volume.uuid}/files/{path} -d '{
"path" : "directory1/directory2/file1" }'
```

== Examples

=== Writing to a new file

The API:

```
POST /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X POST "https://<mgmt-ip>/api/storage/volumes/54c06ce2-5430-11ea-90f9-005056a73aff/files/aNewFile" -H "Content-Type: multipart/form-data"
--form "file=the data to be written to the new file"
```

=== Writing to an existing file

```

# The API:
PATCH /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/54c06ce2-5430-11ea-
90f9-005056a73aff/files/aNewFile?byte_offset=39" -H "Content-Type:
multipart/form-data" --form "file=*here is a little more data"
----

=== Reading a file

----

# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/54c06ce2-5430-11ea-
90f9-005056a73aff/files/aNewFile?byte_offset=0&length=100" -H "Accept:
multipart/form-data"

# Response for file data:
--ec51b3541741ade7
Content-Disposition: form-data; name="bytes_read"
Content-Type: text/plain
66
--ec51b3541741ade7
Content-Disposition: form-data; filename="aNewFile"
Content-Type: application/octet-stream
the data to be written to the new file*here is a little more data
--ec51b3541741ade7--
----

=== Creating a directory

You can use the POST request to create a directory.

----

# The API:
POST /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/dir1" -H 'accept: application/hal+json' -d '{
"type" : "directory", "unix_permissions" : "644" }'

```

```

# The response:
{
  "num_records": 1,
  "records": [
    {
      "path": "dir1",
      "type": "directory",
      "unix_permissions": 644
    }
  ]
}
-----

=== Creating a stream on a file

-----

# The API:
POST /api/storage/volumes/{volume.uuid}/files/{path}?overwrite=true

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes/54c06ce2-5430-11ea-90f9-005056a73aff/files/aNewFile?overwrite=true&byte_offset=1&stream_name=someStream" -H "Content-Type: multipart/form-data" --form "file=the data to be written to the new file"
-----

=== Retrieving the list of files in a directory

-----

# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/d1%2Fd2%2Fd3"

# Response for file records:
{
  "records": [
    {
      "path": "d1/d2/d3",
      "name": ".",
      "type": "directory",
      "_links": {

```



```

    "self": {
      "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3%2F%2E"
    },
    "metadata": {
      "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2F%2E?return_metadata=true"
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "..",
    "type": "directory",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3%2F%2E%2E"
      },
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2F%2E%2E?return_metadata=true"
      }
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "f1",
    "type": "file",
    "_links": {
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2Ffile1?return_metadata=true"
      }
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "d5",
    "type": "directory",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3%2Fd5"
      },
      "metadata": {

```

```

        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-
005056a7d72a/files/d1%2Fd2%2Fd3%2Fd5?return_metadata=true"
    }
}
}
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d1%2Fd2%2Fd3"
  }
}
}
}
-----

```

=== Retrieving a list of files based on file type

You can filter the list of files you retrieve based on multiple file types by including a query parameter in the following format
type="file|symlink"

The API:

```
GET /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/d1%2Fd2%2Fd3?type=file&#124;directory"
```

Response for file records:

```

{
  "records": [
    {
      "path": "d1/d2/d3",
      "name": ".",
      "type": "directory",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-
005056aca658/files/d1%2Fd2%2Fd3%2F%2E"
        },
        "metadata": {
          "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-
005056a7d72a/files/d1%2Fd2%2Fd3%2F%2E?return_metadata=true"
        }
      }
    }
  ]
}

```

```

    }
  },
  {
    "path": "d1/d2/d3",
    "name": "..",
    "type": "directory",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3%2F%2E%2E"
      },
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2F%2E%2E?return_metadata=true"
      }
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "f1",
    "type": "file",
    "_links": {
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2Ffile1?return_metadata=true"
      }
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "d5",
    "type": "directory",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3%2Fd5"
      },
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2Fd5?return_metadata=true"
      }
    }
  }
],
"num_records": 4,

```

```
"_links": {
  "self": {
    "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3"
  }
}
}
```

=== Retrieving the properties of a directory or a file

The API:

```
GET /api/storage/volumes/{volume.uuid}/files/{path}?return_metadata=true
```

The call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/d1%2Fd2%2Fd3%2Ff1?return_metadata=true"
```

Response for file properties:

```
{
  "records": [
    {
      "path": "d1/d2/d3/f1",
      "name": "",
      "type": "file",
      "creation_time": "2019-06-12T21:27:28-04:00",
      "modified_time": "2019-06-12T21:27:28-04:00",
      "changed_time": "2019-06-12T21:27:28-04:00",
      "accessed_time": "2019-06-12T21:27:28-04:00",
      "unix_permissions": 644,
      "owner_id": 54738,
      "group_id": 30,
      "size": 200,
      "hard_links_count": 1,
      "inode_number": 1233,
      "bytes_used": 4096,
      "unique_bytes": 4096,
      "inode_generation": 214488325,
      "is_vm_aligned": false,
      "is_junction": false
    }
  ],
  "num_records": 1,
  "_links": {
```

```
"self": {
  "href": "/api/storage/volumes/da8bb06c-823e-11e9-b790-005056acdc0/files/d1%2Fd2%2Fd3%2Ff1?return_metadata=true"
}
}
}
```

=== Creating a symlink to a relative path

You can use the POST request to create a symlink.

The API:

```
POST /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X POST "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/symlink1" -H 'accept: application/hal+json' -d '{"target" : "d1/f1"}'
```

The response:

```
{
  "num_records": 1,
  "records": [
    {
      "path": "symlink1",
      "target": "d1/f1"
    }
  ]
}
```

=== Retrieving the target of a symlink

You can use the GET request to view the target of a symlink.

The API:

```
GET /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/symlink1?return_metadata=true&fields=target"
```

```
# The response:
{
  "records": [
    {
      "path": "symlink1",
      "target": "d1/f1"
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/storage/volumes/54c06ce2-5430-11ea-90f9-005056a73aff/files/symlink1?return_metadata=true&fields=target"
    }
  }
}
-----
```

=== Retrieving the usage information for a file

You can use the GET request to retrieve the unique bytes held in a file with or without specifying the offset.

The API:

```
GET /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/f1?return_metadata=true&byte_offset=100&length=200"
```

The response:

```
{
  "records": [
    {
      "path": "d1/d2/d3/f1",
      "type": "file",
      "creation_time": "2019-06-12T21:27:28-04:00",
      "modified_time": "2019-06-12T21:27:28-04:00",
      "changed_time": "2019-06-12T21:27:28-04:00",
      "accessed_time": "2019-06-12T21:27:28-04:00",
      "unix_permissions": 644,
      "owner_id": 54738,
      "group_id": 30,
      "size": 200,
    }
  ]
}
```

```

    "hard_links_count": 1,
    "inode_number": 1233,
    "bytes_used": 4096,
    "unique_bytes": 4096,
    "inode_generation": 214488325,
    "is_vm_aligned": false,
    "is_junction": false
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/storage/volumes/cb6b139-8d21-11e9-b926-05056aca658/files/f1?return_metadata=true&byte_offset=100&length=200"
  }
}
}

```

=== Retrieving all information (including analytics) for a directory

The API:

```
GET /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/1ef5d1b2-f9d7-11e9-8043-00505682f860/files/d1?return_metadata=true&fields=**"
```

Response for all fields of the directory:

```

{
"records": [
  {
    "svm": {
      "uuid": "58a996a2-f9d5-11e9-8043-00505682f860",
      "_links": {
        "self": {
          "href": "/api/svm/svms/58a996a2-f9d5-11e9-8043-00505682f860"
        }
      }
    }
  },
  "volume": {
    "uuid": "1ef5d1b2-f9d7-11e9-8043-00505682f860",
    "_links": {
      "self": {

```

```
"href": "/api/storage/volumes/1ef5d1b2-f9d7-11e9-8043-00505682f860"
  }
}
},
"path": "d1",
"type": "directory",
"creation_time": "2019-10-28T23:04:13+00:00",
"modified_time": "2019-10-28T23:10:30+00:00",
"changed_time": "2019-10-28T23:10:30+00:00",
"accessed_time": "2019-10-28T23:10:38+00:00",
"unix_permissions": 755,
"owner_id": 1002,
"group_id": 65533,
"size": 4096,
"hard_links_count": 5,
"inode_number": 96,
"is_empty": false,
"bytes_used": 4096,
"inode_generation": 214514951,
"is_vm_aligned": false,
"is_junction": false,
"analytics": {
  "file_count": 668,
  "bytes_used": 209657856,
  "subdir_count": 18,
  "by_modified_time": {
    "bytes_used": {
      "values": [
        0,
        0,
        0,
        0,
        3112960,
        0,
        14041088,
        20545536,
        0,
        57933824,
        61947904,
        68804608,
        188686336,
        0,
        0,
        0,
        20971520,

```



```
    0
  ],
  "percentages": [
    0,
    0,
    0,
    0,
    1.48,
    0,
    6.7,
    9.8,
    0,
    27.63,
    29.55,
    32.82,
    90,
    0,
    0,
    0,
    10,
    0
  ],
  "labels": [
    "2019-W42",
    "2019-W41",
    "2019-W40",
    "2019-W39",
    "2019-W38",
    "2019-10",
    "2019-09",
    "2019-08",
    "2019-Q4",
    "2019-Q3",
    "2019-Q2",
    "2019-Q1",
    "2019",
    "2018",
    "2017",
    "2016",
    "--2015",
    "unknown"
  ]
}
},
"by_accessed_time": {
  "bytes_used": {
```

```
"values": [  
  102760448,  
  1867776,  
  1245184,  
  2179072,  
  1556480,  
  105873408,  
  9027584,  
  8093696,  
  105873408,  
  23969792,  
  32382976,  
  26460160,  
  188686336,  
  0,  
  0,  
  0,  
  20971520,  
  0  
],  
"percentages": [  
  49.01,  
  0.89,  
  0.59,  
  1.04,  
  0.74,  
  50.5,  
  4.31,  
  3.86,  
  50.5,  
  11.43,  
  15.45,  
  12.62,  
  90,  
  0,  
  0,  
  0,  
  10,  
  0  
],  
"labels": [  
  "2019-W42",  
  "2019-W41",  
  "2019-W40",  
  "2019-W39",  
  "2019-W38",
```

```

        "2019-10",
        "2019-09",
        "2019-08",
        "2019-Q4",
        "2019-Q3",
        "2019-Q2",
        "2019-Q1",
        "2019",
        "2018",
        "2017",
        "2016",
        "--2015",
        "unknown"
    ]
}
}
}
},
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/storage/volumes/1ef5d1b2-f9d7-11e9-8043-00505682f860/files/d1?return_metadata=true&fields=*"
  }
}
}
-----

```

=== Retrieving file system analytics information for a set of histogram buckets

The API:

```
GET /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/d3?type=directory&fields=analytics&analytics.histogram_by_time_labels=2019-Q3,2019-Q2,2019-Q1,2018-Q4"
```

Response with analytics data

```
{
"records": [
```

```

{
  "path": "d3",
  "name": ".",
  "type": "directory",
  "analytics": {
    "file_count": 44,
    "bytes_used": 244240384,
    "subdir_count": 14,
    "by_modified_time": {
      "bytes_used": {
        "values": [
          57344,
          29720576,
          196141056,
          57344
        ],
        "percentages": [
          0.02,
          12.17,
          80.31,
          0.02
        ]
      }
    },
    "by_accessed_time": {
      "bytes_used": {
        "values": [
          69632,
          244170752,
          0,
          0
        ],
        "percentages": [
          0.03,
          99.97,
          0,
          0
        ]
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3%2F%2E"
    }
  }
}

```

```

    "metadata": {
      "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3%2F%2E?return_metadata=true"
    }
  },
  {
    "path": "d3",
    "name": "..",
    "type": "directory",
    "analytics": {
      "file_count": 515,
      "bytes_used": 3034574848,
      "subdir_count": 23,
      "by_modified_time": {
        "bytes_used": {
          "values": [
            61440,
            1756479488,
            214622208,
            1191936
          ],
          "percentages": [
            0,
            57.88,
            7.07,
            0.04
          ]
        }
      },
      "by_accessed_time": {
        "bytes_used": {
          "values": [
            282624,
            3034292224,
            0,
            0
          ],
          "percentages": [
            0.01,
            99.99,
            0,
            0
          ]
        }
      }
    }
  }
}

```

```

    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3%2F%2E%2E"
      },
      "metadata": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3%2F%2E%2E?return_metadata=true"
      }
    }
  },
  {
    "path": "d3",
    "name": "d5",
    "type": "directory",
    "analytics": {
      "file_count": 10,
      "bytes_used": 47648768,
      "subdir_count": 4,
      "by_modified_time": {
        "bytes_used": {
          "values": [
            0,
            29638656,
            0,
            0
          ],
          "percentages": [
            0,
            62.20,
            0,
            0
          ]
        }
      },
      "by_accessed_time": {
        "bytes_used": {
          "values": [
            0,
            47648768,
            0,
            0
          ],
          "percentages": [
            0,

```

```

        100,
        0,
        0
    ]
    }
}
},
"_links": {
    "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3%2Fd5"
    },
    "metadata": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3%2Fd5?return_metadata=true"
    }
}
}
],
"num_records": 3,
"analytics": {
    "by_modified_time": {
        "bytes_used": {
            "labels": [
                "2019-Q3",
                "2019-Q2",
                "2019-Q1",
                "2018-Q4"
            ]
        }
    },
    "by_accessed_time": {
        "bytes_used": {
            "labels": [
                "2019-Q3",
                "2019-Q2",
                "2019-Q1",
                "2018-Q4"
            ]
        }
    }
},
"_links": {
    "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3?type=directory&fields=analytics&analytics.histogram_

```

```
by_time_labels=2019-Q3,2019-Q2,2019-Q1,2018-Q4"
```

```
  }  
}  
}  
----
```

```
=== Identifying the largest subdirectories
```

```
----
```

```
# The API:
```

```
GET /api/storage/volumes/{volume.uuid}/files/{path}
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/1ef5d1b2-f9d7-11e9-  
8043-  
00505682f860/files/d1?fields=analytics.bytes_used&type=directory&order_by=  
analytics.bytes_used%20desc"
```

```
# Response with the largest subdirectories sorted by their usage:
```

```
{  
  "records": [  
    {  
      "path": "d1",  
      "name": "..",  
      "type": "directory",  
      "analytics": {  
        "bytes_used": 56623104  
      }  
    },  
    {  
      "path": "d1",  
      "name": ".",  
      "type": "directory",  
      "analytics": {  
        "bytes_used": 35651584  
      }  
    },  
    {  
      "path": "d1",  
      "name": "biggest",  
      "type": "directory",  
      "analytics": {  
        "bytes_used": 17825792  
      }  
    },  
  ],  
}
```



```

{
  "path": "d1",
  "name": "bigger",
  "type": "directory",
  "analytics": {
    "bytes_used": 10485760
  }
},
{
  "path": "d1",
  "name": "big",
  "type": "directory",
  "analytics": {
    "bytes_used": 5242880
  }
}
],
"num_records": 5,
"_links": {
  "self": {
    "href": "/api/storage/volumes/1ef5d1b2-f9d7-11e9-8043-00505682f860/files/d1?fields=analytics.bytes_used&type=directory&order_by=analytics.bytes_used%20desc"
  }
}
}
}
-----

```

=== Assigning a QoS policy to a file

You can use the PATCH request to assign a QoS policy to a file.

The API:

```
PATCH /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/directory1%2Ffile1" -d '{ "qos_policy": { "name": "policy" } }'
```

The response:

```
{}
```

```
=== Retrieving QoS information for a file
```

You can use the GET request for all fields with `return_metadata="true"` to retrieve QoS information for the file.

```
----
```

```
# The API:
```

```
GET /api/storage/volumes/{volume.uuid}/files/{path}
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/file?return_metadata=true&fields=**"
```

```
# The response:
```

```
{
"records": [
  {
    "svm": {
      "uuid": "42ee3002-67dd-11ea-8508-005056a7b8ac"
    },
    "volume": {
      "uuid": "c05eb66a-685f-11ea-8508-005056a7b8ac"
    },
    "path": "file",
    "type": "lun",
    "creation_time": "2020-03-17T10:58:40-04:00",
    "modified_time": "2020-03-24T18:15:40-04:00",
    "changed_time": "2020-03-24T18:15:40-04:00",
    "accessed_time": "2020-03-24T18:15:40-04:00",
    "unix_permissions": 644,
    "owner_id": 0,
    "group_id": 0,
    "size": 1048576,
    "hard_links_count": 2,
    "inode_number": 96,
    "bytes_used": 1056768,
    "inode_generation": 219748425,
    "is_vm_aligned": false,
    "is_junction": false,
    "is_snapshot": false,
    "qos_policy": {
      "name": "pg1",
      "uuid": "00725264-688f-11ea-8f10-005056a7b8ac"
    }
  }
}
```

```
] ,  
"num_records": 1  
}  
----
```

=== Deleting an entire directory

You can use the DELETE request to remove an entire directory recursively.

The API:

```
DELETE /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/directory1%2Fdirectory2?recurse=true"
```

The response:

```
{  
"job": {  
  "uuid": "27d287e8-fcd4-11e9-b8a4-005056a7b97b",  
  "_links": {  
    "self": {  
      "href": "/api/cluster/jobs/27d287e8-fcd4-11e9-b8a4-005056a7b97b"  
    }  
  }  
}  
}  
}
```

=== Deleting an entire directory with specified throttling threshold

You can specify the maximum number of directory delete operations per second when removing an entire directory recursively.

The API:

```
DELETE /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/directory1%2Fdirectory2?recurse=true&throttle_deletion=100"
```

```
# The response:
{
  "job": {
    "uuid": "27d287e8-fcd4-11e9-b8a4-005056a7b97b",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/27d287e8-fcd4-11e9-b8a4-005056a7b97b"
      }
    }
  }
}
-----
```

=== Deleting an empty directory

You can use the DELETE request to remove an empty directory.

```
# The API:
DELETE /api/storage/volumes/{volume.uuid}/files/{path}
```

```
# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/directory1%2Fdirectory2"
```

```
# The response:
```

```
{}
```

=== Deleting a file

You can use the DELETE request to remove a file.

```
# The API:
DELETE /api/storage/volumes/{volume.uuid}/files/{path}
```

```
# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/directory1%2Ffile2"
```

```
# The response:
```

```
{}
```

=== Renaming a file

You can use the PATCH request to rename a file.

The API:

```
PATCH /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/directory1%2Fdirectory2%2Ffile1" -d '{"path": "directory1/file2" }'
```

The response:

```
{}
```

=== File truncating

You can use the PATCH request to change the size of a file.

The API:

```
PATCH /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/abc.txt" -d '{"size":100 }'
```

The response:

```
{}
```

=== Renaming a directory

You can use the PATCH request to rename a directory.

The API:

```
PATCH /api/storage/volumes/{volume.uuid}/files/{path}
```

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/directory1%2Fdirectory2" -d '{"path": "d3/d4" }'
```

```
# The response:
```

```
{}
```

```
----
```

```
[[IDc2dc63dc060250d5e83bdac0e4bdd4cd]]
```

```
= Delete an existing file or directory
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-  
block]#`/storage/volumes/{volume.uuid}/files/{path}`#
```

```
*Introduced In:* 9.8
```

Deletes an existing file or directory. Query-based DELETE operations are not supported.

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|volume.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Volume UUID
```

```
|path
```

```
|string
```

```
|path
```

```
|True
```

a|The relative path of a directory in the volume. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided.

```
|recurse
|boolean
|query
|False
a|Delete an entire directory. The behaviour of this call is equivalent to
rm -rf.
```

* Default value:

```
|throttle_deletion
|integer
|query
|False
a|The maximum number of directory delete operations per second. A valid
throttle_deletion number is an integer from 10 to 100000.
```

```
|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.
```

* Default value:

```
|return_timeout
|integer
|query
|False
a|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

== Response

Status: 202, Accepted

== Error

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
------------	-------------

131074	No such file or directory.
--------	----------------------------

131102	Read-only file system.
--------	------------------------

131138	Directory not empty.
--------	----------------------

918235	A volume with UUID {volume.uuid} was not found.
--------	---

6488081	The \{field} field is not supported for DELETE operations.
---------	--

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===


```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID8ae6db6bd0a723e21f31d9543e4c6a45]]
= Retrieve files and directories

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/volumes/{volume.uuid}/files/{path}`#

```

Introduced In: 9.7

Retrieves a list of files and directories for a given directory or returns only the properties of a single given directory or file of a 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 property. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

*** `analytics`

*** `qos_policy.name`

*** `qos_policy.uuid`

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|volume.uuid

|string

|path

|True

a|Volume UUID

|path

|string

|path

|True

a|Relative path of a file or directory in the volume. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided.

```
|byte_offset
|integer
|query
|False
a|The file offset to start reading from.
```

* Introduced in: 9.8

```
|length
|integer
|query
|False
a|Length of the range in bytes.
```

* Introduced in: 9.8

```
|return_metadata
|boolean
|query
|False
a|If true, the request returns metadata for the the directory or file
specified in the path.
```

* Introduced in: 9.8

* Default value:

```
|group_id
|integer
|query
|False
a|Filter by group_id
```

```
|owner_id
|integer
|query
|False
a|Filter by owner_id
```

```
|unique_bytes
|integer
|query
|False
```

a|Filter by unique_bytes

* Introduced in: 9.8

|fill_enabled

|boolean

|query

|False

a|Filter by fill_enabled

* Introduced in: 9.8

|volume.name

|string

|query

|False

a|Filter by volume.name

|creation_time

|string

|query

|False

a|Filter by creation_time

|constituent.uuid

|string

|query

|False

a|Filter by constituent.uuid

* Introduced in: 9.10

|constituent.name

|string

|query

|False

a|Filter by constituent.name

* Introduced in: 9.10

|size

```
|integer  
|query  
|False  
a|Filter by size
```

```
|changed_time  
|string  
|query  
|False  
a|Filter by changed_time
```

```
|target  
|string  
|query  
|False  
a|Filter by target
```

* Introduced in: 9.8

```
|is_snapshot  
|boolean  
|query  
|False  
a|Filter by is_snapshot
```

* Introduced in: 9.8

```
|is_junction  
|boolean  
|query  
|False  
a|Filter by is_junction
```

```
|name  
|string  
|query  
|False  
a|Filter by name
```

```
|inode_generation  
|integer
```

```
|query
|False
a|Filter by inode_generation
```

```
|qos_policy.uuid
|string
|query
|False
a|Filter by qos_policy.uuid
```

* Introduced in: 9.8

```
|qos_policy.name
|string
|query
|False
a|Filter by qos_policy.name
```

* Introduced in: 9.8

```
|modified_time
|string
|query
|False
a|Filter by modified_time
```

```
|is_vm_aligned
|boolean
|query
|False
a|Filter by is_vm_aligned
```

```
|hard_links_count
|integer
|query
|False
a|Filter by hard_links_count
```

```
|analytics.by_modified_time.bytes_used.newest_label
|string
|query
```

```
|False
a|Filter by analytics.by_modified_time.bytes_used.newest_label

* Introduced in: 9.8

|analytics.by_modified_time.bytes_used.oldest_label
|string
|query
|False
a|Filter by analytics.by_modified_time.bytes_used.oldest_label

* Introduced in: 9.8

|analytics.by_modified_time.bytes_used.percentages
|number
|query
|False
a|Filter by analytics.by_modified_time.bytes_used.percentages

* Introduced in: 9.8

|analytics.by_modified_time.bytes_used.labels
|string
|query
|False
a|Filter by analytics.by_modified_time.bytes_used.labels

* Introduced in: 9.8

|analytics.by_modified_time.bytes_used.values
|integer
|query
|False
a|Filter by analytics.by_modified_time.bytes_used.values

* Introduced in: 9.8

|analytics.bytes_used
|integer
|query
|False
a|Filter by analytics.bytes_used
```


* Introduced in: 9.8

```
|analytics.incomplete_data  
|boolean  
|query  
|False  
a|Filter by analytics.incomplete_data
```

* Introduced in: 9.12

```
|analytics.subdir_count  
|integer  
|query  
|False  
a|Filter by analytics.subdir_count
```

* Introduced in: 9.8

```
|analytics.file_count  
|integer  
|query  
|False  
a|Filter by analytics.file_count
```

* Introduced in: 9.8

```
|analytics.by_accessed_time.bytes_used.newest_label  
|string  
|query  
|False  
a|Filter by analytics.by_accessed_time.bytes_used.newest_label
```

* Introduced in: 9.8

```
|analytics.by_accessed_time.bytes_used.oldest_label  
|string  
|query  
|False  
a|Filter by analytics.by_accessed_time.bytes_used.oldest_label
```

* Introduced in: 9.8

```
|analytics.by_accessed_time.bytes_used.percentages
|number
|query
|False
a|Filter by analytics.by_accessed_time.bytes_used.percentages
```

* Introduced in: 9.8

```
|analytics.by_accessed_time.bytes_used.labels
|string
|query
|False
a|Filter by analytics.by_accessed_time.bytes_used.labels
```

* Introduced in: 9.8

```
|analytics.by_accessed_time.bytes_used.values
|integer
|query
|False
a|Filter by analytics.by_accessed_time.bytes_used.values
```

* Introduced in: 9.8

```
|overwrite_enabled
|boolean
|query
|False
a|Filter by overwrite_enabled
```

* Introduced in: 9.8

```
|inode_number
|integer
|query
|False
a|Filter by inode_number
```

```
|unix_permissions
|integer
```

```
|query
|False
a|Filter by unix_permissions
```

```
|bytes_used
|integer
|query
|False
a|Filter by bytes_used
```

```
|is_empty
|boolean
|query
|False
a|Filter by is_empty
```

```
|type
|string
|query
|False
a|Filter by type
```

```
|accessed_time
|string
|query
|False
a|Filter by accessed_time
```

```
|analytics.histogram_by_time_labels
|array[string]
|query
|False
a|Request that returned
xref:{relative_path}analytics_histogram_by_time.html[analytics_histogram_b
y_time] objects including values associated with the specified labels.
```

As described in the object description, the labels may take the following forms:

- `_partial-date_`
- `+++<tt>+++--+++</tt>+++ _partial-date_`
- `_partial-date_ +++<tt>+++--+++</tt>+++`
- `_partial-date_ +++<tt>+++--+++</tt>+++ _partial-date_`

+++<tt>+++unknown+++</tt>+++

Intervals that the system would not normally return may be specified. In this case, the appropriate values and percentages summarizing all files with a time-based attribute within the indicated period of time are calculated and returned in the response. However, there are some restrictions:

Any `__partial-date__` specified as the beginning or end of an interval must be tracked by the system. Valid `__partial-date__`s may be determined by making an OPTIONS request to the

+++<tt>+++/storage/volumes/{volume.uuid}/files/{path}+++</tt>+++ endpoint.

Intervals may not mix week-based `__partial-date__`s in the form `__yyyy__-W__ww__` with other types of `__partial-date__`s.

* Introduced in: 9.8

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|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
a|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
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|analytics
|link:#analytics[analytics]
a|Additional file system analytics information that is invariant amongst
all elements in the collection.

This property is only populated if file system analytics is enabled on the
containing volume.

This analytics object captures properties that are invariant amongst all
elements included in the `records` array. The invariant properties are
included here, rather than within the information for each element, to
avoid returning an excessive amount of duplicated information when the

```

collection is large.

```
|num_records  
|integer  
a|Number of records.
```

```
|records  
|array[link:#file_info[file_info]]  
a|
```

```
|===
```

.Example response
[%collapsible%closed]

```
=====
```

```
[source,json,subs=+macros]
```

```
{  
  "_links": {  
    "next": {  
      "href": "/api/resourcelink"  
    },  
    "self": {  
      "href": "/api/resourcelink"  
    }  
  },  
  "analytics": {  
    "by_accessed_time": {  
      "bytes_used": {  
        "labels": [  
          "2019-07",  
          "2019-06",  
          "2019-05",  
          "2019",  
          "2018",  
          "--2017",  
          "unknown"  
        ]  
      }  
    },  
    "by_modified_time": {  
      "bytes_used": {  
        "labels": [  
          "2019-07",
```

```
        "2019-06",
        "2019-05",
        "2019",
        "2018",
        "--2017",
        "unknown"
    ]
}
},
"num_records": 1,
"records": {
  "_links": {
    "metadata": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "accessed_time": "2019-06-12 15:00:16 +0000",
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": "2019-07",
        "oldest_label": "2019-07",
        "percentages": [
          0.1,
          11.24,
          0.18,
          15.75,
          0.75,
          83.5,
          0
        ],
        "values": [
          15925248,
```

```
        1735569408,  
        27672576,  
        2430595072,  
        116105216,  
        12889948160,  
        0  
    ]  
    }  
},  
"by_modified_time": {  
    "bytes_used": {  
        "labels": [  
            "2019-07",  
            "2019-06",  
            "2019-05",  
            "2019",  
            "2018",  
            "--2017",  
            "unknown"  
        ],  
        "newest_label": "2019-07",  
        "oldest_label": "2019-07",  
        "percentages": [  
            0.1,  
            11.24,  
            0.18,  
            15.75,  
            0.75,  
            83.5,  
            0  
        ],  
        "values": [  
            15925248,  
            1735569408,  
            27672576,  
            2430595072,  
            116105216,  
            12889948160,  
            0  
        ]  
    }  
},  
"bytes_used": 15436648448,  
"file_count": 21134,  
"subdir_count": 35  
},
```



```
"bytes_used": 4096,
"changed_time": "2019-06-12 15:00:16 +0000",
"constituent": {
  "name": "fg_0001",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"creation_time": "2019-06-12 15:00:16 +0000",
"group_id": 30,
"hard_links_count": 1,
"inode_generation": 214753547,
"inode_number": 1695,
"is_empty": "",
"is_junction": "",
"is_snapshot": "",
"is_vm_aligned": "",
"modified_time": "2019-06-12 15:00:16 +0000",
"name": "test_file",
"owner_id": 54738,
"path": "d1/d2/d3",
"qos_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "qos1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": 200,
"target": "some_directory/some_other_directory/some_file",
"type": "file",
"unique_bytes": 4096,
"unix_permissions": 493,
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
}
====
```

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```
|===  
|Name  
|Type  
|Description
```

```
|href  
|string  
a|
```

```
|===
```

```
[#_links]  
[.api-collapsible-fifth-title]  
_links
```

```
[cols=3*,options=header]
```

```
|===  
|Name  
|Type  
|Description
```

```
|next  
|link:href[href]  
a|
```

```
|self  
|link:href[href]  
a|
```

```
|===
```

```
[#bytes_used]  
[.api-collapsible-fifth-title]  
bytes_used
```

Number of bytes used on-disk, broken down by date of last access.

```
[cols=3*,options=header]
```

```
|===  
|Name  
|Type  
|Description
```

```
|labels
```

```
|array[string]
a|Labels for this histogram
```

```
|===
```

```
[#by_accessed_time]
[.api-collapsible-fifth-title]
by_accessed_time
```

File system analytics information, broken down by date of last access.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|bytes_used
|link:#bytes_used[bytes_used]
```

a|Number of bytes used on-disk, broken down by date of last access.

```
|===
```

```
[#bytes_used]
[.api-collapsible-fifth-title]
bytes_used
```

Number of bytes used on-disk, broken down by date of last modification.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|labels
|array[string]
a|Labels for this histogram
```

```
|===
```

```
[#by_modified_time]
[.api-collapsible-fifth-title]
by_modified_time
```

File system analytics information, broken down by date of last modification.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|bytes_used
```

```
|link:#bytes_used[bytes_used]
```

```
a|Number of bytes used on-disk, broken down by date of last modification.
```

```
|===
```

```
[#analytics]
```

```
[.api-collapsible-fifth-title]
```

```
analytics
```

Additional file system analytics information that is invariant amongst all elements in the collection.

This property is only populated if file system analytics is enabled on the containing volume.

This analytics object captures properties that are invariant amongst all elements included in the `records` array. The invariant properties are included here, rather than within the information for each element, to avoid returning an excessive amount of duplicated information when the collection is large.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|by_accessed_time
|link:#by_accessed_time[by_accessed_time]
a|File system analytics information, broken down by date of last access.
```

```
|by_modified_time
|link:#by_modified_time[by_modified_time]
a|File system analytics information, broken down by date of last
modification.
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|metadata
|link:#href[href]
a|
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#bytes_used]
[.api-collapsible-fifth-title]
bytes_used
```

Number of bytes used on-disk, broken down by date of last access.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|labels
|array[string]
a|Labels for this histogram
```

```
|newest_label
|string
```

a|Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:
a partial date in an extended ISO8601 representation
an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval
the string literal "unknown"

For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00--2017-12-31T23:59:59. Similarly, the interval "2018-05--2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00--2018-07-31T23:59:59.

The following extensions to ISO8601 are used:

Quarters may be specified. The form `__yyyy_-Q__q__` is used to represent the `__q__`th quarter of the year `__yyyy_`. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00--2019-06-30T23:59:59.

Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval includes points in time through the end of the current week.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

```
|oldest_label
|string
```

a|Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:

a partial date in an extended ISO8601 representation
an interval between partial dates in an extended ISO8601 representation,
where "--" is used to separate the beginning and end of the interval
the string literal "unknown"

For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00--2017-12-31T23:59:59. Similarly, the interval "2018-05--2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00--2018-07-31T23:59:59.

The following extensions to ISO8601 are used:

Quarters may be specified. The form `__yyyy_-Q__q__` is used to represent the `__q__`th quarter of the year `__yyyy_`. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00--2019-06-30T23:59:59.

Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval includes points in time through the end of the current week.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

```
|percentages  
|array[number]  
a|Percentages for this histogram
```

```
|values  
|array[integer]  
a|Values for this histogram
```

```
|===
```

```
[#bytes_used]  
[.api-collapsible-fifth-title]
```


bytes_used

Number of bytes used on-disk, broken down by date of last modification.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|labels
```

```
|array[string]
```

```
a|Labels for this histogram
```

```
|newest_label
```

```
|string
```

```
a|Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:<ul>  
  a partial date in an extended ISO8601 representation  
  an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval  
  the string literal "unknown"
```

```
</ul>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00--2017-12-31T23:59:59. Similarly, the interval "2018-05--2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00--2018-07-31T23:59:59.
```

```
The following extensions to ISO8601 are used:<ul>
```

```
  Quarters may be specified. The form yyyy_Q_q is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00--2019-06-30T23:59:59.
```

```
  Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval includes points in time through the end of the current week.
```

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

|oldest_label

|string

a|Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:

- a partial date in an extended ISO8601 representation
- an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval
- the string literal "unknown"

For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00--2017-12-31T23:59:59. Similarly, the interval "2018-05--2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00--2018-07-31T23:59:59.

The following extensions to ISO8601 are used:

Quarters may be specified. The form `_yyyy_-Q__q__` is used to represent the `__q__`th quarter of the year `_yyyy_`. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00--2019-06-30T23:59:59.

Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval includes points in time through the end of the current week.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

|percentages

|array[number]

a|Percentages for this histogram

|values

```
|array[integer]
a|Values for this histogram
```

```
|===
```

```
[#analytics]
[.api-collapsible-fifth-title]
analytics
```

Additional file system analytics information summarizing all descendents of a directory.

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a `xref:{relative_path}file-info-response(#model-file-info-response)`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the `file-info-response(#model-file-info-response)`. This avoids an excessive amount of duplicated information when a `get-storage-volumes-files-.html<<model-file_info_response,file_info_response>>`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the `<<model-file_info_response,file_info_response>>`. This avoids an excessive amount of duplicated information when a `[GET /storage/volumes/{volume.uuid}/files/{path}]` call returns a large collection.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|by_accessed_time
|link:#by_accessed_time[by_accessed_time]
a|File system analytics information, broken down by date of last access.
```

```
|by_modified_time
|link:#by_modified_time[by_modified_time]
a|File system analytics information, broken down by date of last
modification.
```

```
|bytes_used
|integer
a|Number of bytes used on-disk
```

```
|file_count
|integer
a|Number of descendants
```

```
|incomplete_data
|boolean
a|Returns true if data collection is incomplete for this directory tree.
```

```
|subdir_count
|integer
a|Number of sub directories
```

```
|===
```

```
[#constituent]
[.api-collapsible-fifth-title]
constituent
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|FlexGroup volume constituent name.
```

```
|uuid
|string
```

```
a|FlexGroup volume constituent UUID.
```

```
|===
```

```
[#_links]  
[.api-collapsible-fifth-title]  
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|self  
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#qos_policy]  
[.api-collapsible-fifth-title]  
qos_policy
```

The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both ``qos_policy.uuid`` and ``qos_policy.name`` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property ``qos_policy.name`` in a PATCH request to an empty string `""` or `"none"`.

NOTE: Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on `/storage/luns` to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the QoS policy. To remove the file from a QoS policy, set this property to an empty string "" or set it to "none" in a PATCH request.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the QoS policy. Valid in PATCH.
```

```
|===
```

```
[#volume]
```

```
[.api-collapsible-fifth-title]
```

```
volume
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the volume.
```

```
|uuid
```

```
|string
```

```
a|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
```

```
|===
```

```
[#file_info]  
[.api-collapsible-fifth-title]  
file_info
```

Information about a single file.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|accessed_time  
|string  
a|Last access time of the file in date-time format.
```

```
|analytics  
|link:#analytics[analytics]  
a|Additional file system analytics information summarizing all descendents  
of a directory.
```

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a `xref:{relative_path}file-info-response(#model-file-info-response)`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the `file-info-response(#model-file-info-response)`. This avoids an excessive amount of duplicated information when a `get-storage-volumes-files-.html<<model-file_info_response,file_info_response>>`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the

collection. The invariant information is instead available via the analytics property of the <<model-file_info_response,file_info_response>>. This avoids an excessive amount of duplicated information when a [GET /storage/volumes/{volume.uuid}/files/{path}] call returns a large collection.

|bytes_used

|integer

a|The actual number of bytes used on disk by this file. If byte_offset and length parameters are specified, this will return the bytes used by the file within the given range.

|changed_time

|string

a|Last time data or attributes changed on the file in date-time format.

|constituent

|link:#constituent[constituent]

a|

|creation_time

|string

a|Creation time of the file in date-time format.

|fill_enabled

|boolean

a|Returns "true" if the space reservation is enabled. The field overwrite_enabled must also be set to the same value as this field.

|group_id

|integer

a|The integer ID of the group of the file owner.

|hard_links_count

|integer

a|The number of hard links to the file.

|inode_generation

|integer

a|Inode generation number.

|inode_number
|integer
a|The file inode number.

|is_empty
|boolean
a|Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set.

|is_junction
|boolean
a|Returns "true" if the directory is a junction.

|is_snapshot
|boolean
a|Returns "true" if the directory is a Snapshot copy.

|is_vm_aligned
|boolean
a|Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume.

|modified_time
|string
a|Last data modification time of the file in date-time format.

|name
|string
a|Name of the file.

|overwrite_enabled
|boolean
a|Returns "true" if the space reservation for overwrites is enabled. The

field `fill_enabled` must also be set to the same value as this field.

|owner_id

|integer

a|The integer ID of the file owner.

|path

|string

a|Path of the file.

|qos_policy

|link:#qos_policy[qos_policy]

a|The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both ``qos_policy.uuid`` and ``qos_policy.name`` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property ``qos_policy.name`` in a PATCH request to an empty string `""` or `"none"`.

NOTE: Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on `/storage/luns` to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

|size

|integer

a|The size of the file, in bytes.

|target

|string

a|The relative or absolute path contained in a symlink, in the form `+++<some>+++ /+++<path>+++ .+++</path>++++++</some>+++`

|type

|string

a|Type of the file.

|unique_bytes

|integer

a|Number of bytes uniquely held by this file. If `byte_offset` and `length`

parameters are specified, this will return bytes uniquely held by the file within the given range.

|unix_permissions

|integer

a|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). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group.

|volume

|link:#volume[volume]

a|

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID07d3127430c05e357a1b2526d9a6a229]]
= Write to an existing file with the supplied data

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/volumes/{volume.uuid}/files/{path}`#

*Introduced In:* 9.8

Writes to an existing file with the supplied data or modifies the size,
name, space reservation information, QoS policy, or hole range information
of a file. Query-based PATCH operations are not supported.

```

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|volume.uuid

|string

|path

|True

a|Volume UUID

|path

|string

|path

|True

a|Relative path of a file in the volume. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided.

|byte_offset

|integer

|query

|False

a|How many bytes into the file to begin writing. Use -1 to append (default).

|overwrite

|boolean

|query

|False

a|If false, and the file exists, the write will fail. Default is false.

|stream_name

|string

|query

|False

a|Name of stream associated with the file to write data to.

```
|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.
```

* Default value:

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
 |_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|accessed_time
```

```
|string
```

```
a|Last access time of the file in date-time format.
```

```
|analytics
```

```
|link:#analytics[analytics]
```

```
a|Additional file system analytics information summarizing all descendents of a directory.
```

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a `xref:{relative_path}file-info-response(#model-file-info-response)`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the `file-info-response(#model-file-info-response)`. This avoids an excessive amount of duplicated information when a get-

storage-volumes-files-.html<<model-file_info_response,file_info_response>>, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the <<model-file_info_response,file_info_response>>. This avoids an excessive amount of duplicated information when a [GET /storage/volumes/{volume.uuid}/files/{path}] call returns a large collection.

|bytes_used

|integer

a|The actual number of bytes used on disk by this file. If byte_offset and length parameters are specified, this will return the bytes used by the file within the given range.

|changed_time

|string

a|Last time data or attributes changed on the file in date-time format.

|constituent

|link:#constituent[constituent]

a|

|creation_time

|string

a|Creation time of the file in date-time format.

|fill_enabled

|boolean

a|Returns "true" if the space reservation is enabled. The field overwrite_enabled must also be set to the same value as this field.

|group_id

|integer

a|The integer ID of the group of the file owner.

|hard_links_count

|integer

a|The number of hard links to the file.

|inode_generation
|integer
a|Inode generation number.

|inode_number
|integer
a|The file inode number.

|is_empty
|boolean
a|Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set.

|is_junction
|boolean
a|Returns "true" if the directory is a junction.

|is_snapshot
|boolean
a|Returns "true" if the directory is a Snapshot copy.

|is_vm_aligned
|boolean
a|Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume.

|modified_time
|string
a|Last data modification time of the file in date-time format.

|name
|string
a|Name of the file.

|overwrite_enabled

|boolean

a|Returns "true" if the space reservation for overwrites is enabled. The field fill_enabled must also be set to the same value as this field.

|owner_id

|integer

a|The integer ID of the file owner.

|path

|string

a|Path of the file.

|qos_policy

|link:#qos_policy[qos_policy]

a|The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".

NOTE: Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

|size

|integer

a|The size of the file, in bytes.

|target

|string

a|The relative or absolute path contained in a symlink, in the form +++<some>+++ /+++<path>+++ .+++</path>++++++</some>+++

|type

|string

a|Type of the file.

|unique_bytes

|integer

a|Number of bytes uniquely held by this file. If byte_offset and length parameters are specified, this will return bytes uniquely held by the file within the given range.

|unix_permissions

|integer

a|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). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group.

|volume

|link:#volume[volume]

a|

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "metadata": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "accessed_time": "2019-06-12 15:00:16 +0000",
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "labels": [
          "2019-07",
          "2019-06",

```

```

    "2019-05",
    "2019",
    "2018",
    "--2017",
    "unknown"
  ],
  "newest_label": "2019-07",
  "oldest_label": "2019-07",
  "percentages": [
    0.1,
    11.24,
    0.18,
    15.75,
    0.75,
    83.5,
    0
  ],
  "values": [
    15925248,
    1735569408,
    27672576,
    2430595072,
    116105216,
    12889948160,
    0
  ]
}
},
"by_modified_time": {
  "bytes_used": {
    "labels": [
      "2019-07",
      "2019-06",
      "2019-05",
      "2019",
      "2018",
      "--2017",
      "unknown"
    ],
    "newest_label": "2019-07",
    "oldest_label": "2019-07",
    "percentages": [
      0.1,
      11.24,
      0.18,
      15.75,

```

```

    0.75,
    83.5,
    0
  ],
  "values": [
    15925248,
    1735569408,
    27672576,
    2430595072,
    116105216,
    12889948160,
    0
  ]
}
},
"bytes_used": 15436648448,
"file_count": 21134,
"subdir_count": 35
},
"bytes_used": 4096,
"changed_time": "2019-06-12 15:00:16 +0000",
"constituent": {
  "name": "fg__0001",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"creation_time": "2019-06-12 15:00:16 +0000",
"group_id": 30,
"hard_links_count": 1,
"inode_generation": 214753547,
"inode_number": 1695,
"is_empty": "",
"is_junction": "",
"is_snapshot": "",
"is_vm_aligned": "",
"modified_time": "2019-06-12 15:00:16 +0000",
"name": "test_file",
"owner_id": 54738,
"path": "d1/d2/d3",
"qos_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"name": "qos1",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"

```

```
},
"size": 200,
"target": "some_directory/some_other_directory/some_file",
"type": "file",
"unique_bytes": 4096,
"unix_permissions": 493,
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"name": "volume1",
"uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
| Error Code | Description
| 918235
| A volume with UUID {volume.uuid} was not found.
| 6488081
| The \{field} field is not supported for PATCH operations.
| 6488082
| Failed to rename \{path}.
| 6488083
| Failed to rename \{path} to \{path} because a directory named \{path}
already exists.
|===
```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name

```

```
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|metadata
|link:href[href]
a|
```

```
|self
|link:href[href]
a|
```

```
|===
```

```
[#bytes_used]
[.api-collapsible-fifth-title]
bytes_used
```

Number of bytes used on-disk, broken down by date of last access.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|labels
|array[string]
a|Labels for this histogram
```

|newest_label

|string

a|Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:
a partial date in an extended ISO8601 representation
an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval
the string literal "unknown"

For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00--2017-12-31T23:59:59. Similarly, the interval "2018-05--2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00--2018-07-31T23:59:59.

The following extensions to ISO8601 are used:

Quarters may be specified. The form `__yyyy_-Q__q__` is used to represent the `__q__`th quarter of the year `__yyyy_`. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00--2019-06-30T23:59:59.

Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval includes points in time through the end of the current week.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

|oldest_label

|string

a|Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:
a partial date in an extended ISO8601 representation
an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval
the string literal "unknown"

For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00--2017-12-31T23:59:59. Similarly, the interval "2018-05--2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00--2018-07-31T23:59:59.

The following extensions to ISO8601 are used:

Quarters may be specified. The form `_yyyy_-Q__q__` is used to represent the `__q__`th quarter of the year `_yyyy_`. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00--2019-06-30T23:59:59.

Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval includes points in time through the end of the current week.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

```
|percentages
|array[number]
a|Percentages for this histogram
```

```
|values
|array[integer]
a|Values for this histogram
```

```
|===
```

```
[#by_accessed_time]
[.api-collapsible-fifth-title]
by_accessed_time
```

File system analytics information, broken down by date of last access.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|bytes_used
|link:#bytes_used[bytes_used]
a|Number of bytes used on-disk, broken down by date of last access.

|===

[#bytes_used]
[.api-collapsible-fifth-title]
bytes_used

Number of bytes used on-disk, broken down by date of last modification.
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|labels
|array[string]
a|Labels for this histogram

|newest_label
|string
a|Each label indicates the period of time the corresponding data is
associated with. A label can take one of the following forms:<ul>
  a partial date in an extended ISO8601 representation
  an interval between partial dates in an extended ISO8601 representation,
where "--" is used to separate the beginning and end of the interval
  the string literal "unknown"

</ul>For partial dates and partial date intervals where components of a
date are unspecified, the label allows for any valid normalized values the
unspecified components might take. For example, the label "2017" allows
for any time within the year 2017. Essentially, this is the fully
specified interval 2017-01-01T00:00:00--2017-12-31T23:59:59. Similarly,
the interval "2018-05--2018-07" allows for any time within the months of
```

May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00--2018-07-31T23:59:59.

The following extensions to ISO8601 are used:

Quarters may be specified. The form `__yyyy_-Q__q__` is used to represent the `__q__`th quarter of the year `__yyyy_`. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00--2019-06-30T23:59:59.

Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval includes points in time through the end of the current week.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

|oldest_label

|string

a|Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:

- a partial date in an extended ISO8601 representation
- an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval
- the string literal "unknown"

For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00--2017-12-31T23:59:59. Similarly, the interval "2018-05--2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00--2018-07-31T23:59:59.

The following extensions to ISO8601 are used:

Quarters may be specified. The form `__yyyy_-Q__q__` is used to represent the `__q__`th quarter of the year `__yyyy_`. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00--

2019-06-30T23:59:59.

Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval includes points in time through the end of the current week.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

```
|percentages
|array[number]
a|Percentages for this histogram
```

```
|values
|array[integer]
a|Values for this histogram
```

```
|===
```

```
[#by_modified_time]
[.api-collapsible-fifth-title]
by_modified_time
```

File system analytics information, broken down by date of last modification.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|bytes_used
|link:#bytes_used[bytes_used]
a|Number of bytes used on-disk, broken down by date of last modification.
```

```
|===
```

```
[#analytics]
```

```
[.api-collapsible-fifth-title]
```

```
analytics
```

Additional file system analytics information summarizing all descendents of a directory.

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a `xref:{relative_path}file-info-response(#model-file-info-response)`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the `file-info-response(#model-file-info-response)`. This avoids an excessive amount of duplicated information when a `get-storage-volumes-files-.html<<model-file_info_response,file_info_response>>`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the `<<model-file_info_response,file_info_response>>`. This avoids an excessive amount of duplicated information when a `[GET /storage/volumes/{volume.uuid}/files/{path}]` call returns a large collection.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|by_accessed_time
```

```
|link:#by_accessed_time[by_accessed_time]
```

```
a|File system analytics information, broken down by date of last access.
```

```
|by_modified_time
```

```
|link:#by_modified_time[by_modified_time]
```

```
a|File system analytics information, broken down by date of last modification.
```

```
|bytes_used
```

```

|integer
a|Number of bytes used on-disk

|file_count
|integer
a|Number of descendants

|incomplete_data
|boolean
a|Returns true if data collection is incomplete for this directory tree.

|subdir_count
|integer
a|Number of sub directories

|===

[#constituent]
[.api-collapsible-fifth-title]
constituent

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|FlexGroup volume constituent name.

|uuid
|string
a|FlexGroup volume constituent UUID.

|===

[#_links]
[.api-collapsible-fifth-title]

```

`_links`

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#qos_policy]
```

```
[.api-collapsible-fifth-title]
```

`qos_policy`

The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both ``qos_policy.uuid`` and ``qos_policy.name`` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property ``qos_policy.name`` in a PATCH request to an empty string `""` or `"none"`.

NOTE: Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on `/storage/luns` to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the QoS policy. To remove the file from a QoS policy, set this property to an empty string "" or set it to "none" in a PATCH request.
```

```
|uuid
|string
a|The unique identifier of the QoS policy. Valid in PATCH.
```

```
|===
```

```
[#volume]
[.api-collapsible-fifth-title]
volume
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the volume.
```

```
|uuid
|string
a|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
```

```
|===
```

```
[#file_info]
[.api-collapsible-fifth-title]
file_info
```


Information about a single file.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|accessed_time
```

```
|string
```

```
a|Last access time of the file in date-time format.
```

```
|analytics
```

```
|link:#analytics[analytics]
```

```
a|Additional file system analytics information summarizing all descendents of a directory.
```

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a `xref:{relative_path}file-info-response(#model-file-info-response)`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the `file-info-response(#model-file-info-response)`. This avoids an excessive amount of duplicated information when a `get-storage-volumes-files-.html<<model-file_info_response,file_info_response>>`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the `<<model-file_info_response,file_info_response>>`. This avoids an excessive amount of duplicated information when a `[GET /storage/volumes/{volume.uuid}/files/{path}]` call returns a large collection.

```
|bytes_used
```

```
|integer
a|The actual number of bytes used on disk by this file. If byte_offset and
length parameters are specified, this will return the bytes used by the
file within the given range.

|changed_time
|string
a|Last time data or attributes changed on the file in date-time format.

|constituent
|link:#constituent[constituent]
a|

|creation_time
|string
a|Creation time of the file in date-time format.

|fill_enabled
|boolean
a|Returns "true" if the space reservation is enabled. The field
overwrite_enabled must also be set to the same value as this field.

|group_id
|integer
a|The integer ID of the group of the file owner.

|hard_links_count
|integer
a|The number of hard links to the file.

|inode_generation
|integer
a|Inode generation number.

|inode_number
|integer
a|The file inode number.

|is_empty
```

|boolean

a|Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set.

|is_junction

|boolean

a|Returns "true" if the directory is a junction.

|is_snapshot

|boolean

a|Returns "true" if the directory is a Snapshot copy.

|is_vm_aligned

|boolean

a|Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume.

|modified_time

|string

a|Last data modification time of the file in date-time format.

|name

|string

a|Name of the file.

|overwrite_enabled

|boolean

a|Returns "true" if the space reservation for overwrites is enabled. The field fill_enabled must also be set to the same value as this field.

|owner_id

|integer

a|The integer ID of the file owner.

|path
|string
a|Path of the file.

|qos_policy
|link:#qos_policy[qos_policy]
a|The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".

NOTE: Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

|size
|integer
a|The size of the file, in bytes.

|target
|string
a|The relative or absolute path contained in a symlink, in the form +++<some>+++ /+++<path>+++ .+++</path>++++++</some>+++

|type
|string
a|Type of the file.

|unique_bytes
|integer
a|Number of bytes uniquely held by this file. If byte_offset and length parameters are specified, this will return bytes uniquely held by the file within the given range.

|unix_permissions
|integer
a|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). The

first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group.

```
|volume  
|link:#volume[volume]  
a|
```

```
|===
```

```
[#error_arguments]  
[.api-collapsible-fifth-title]  
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|code  
|string  
a|Argument code
```

```
|message  
|string  
a|Message argument
```

```
|===
```

```
[#error]  
[.api-collapsible-fifth-title]  
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[ID551ae6aa0ba96729ab9a594cdb5e699f]]
```

```
= Create a new file with the supplied data
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-  
block]#`/storage/volumes/{volume.uuid}/files/{path}`#
```

```
*Introduced In:* 9.8
```

Creates a new file with the supplied data, a new directory or a new symlink.

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume UUID

|path
|string
|path
|True
a|Relative path of a new file, directory or symlink. The path field
requires using "%2E" to represent "." and "%2F" to represent "/" for the
path provided.

|byte_offset
|integer
|query
|False
a|How many bytes into the file to begin writing. Use -1 to append
(default).

|overwrite
|boolean
|query
|False
a|If false, and the file exists, the write will fail. Default is false.

|stream_name
|string
|query
|False
a|Name of stream associated with the file to write data to.

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.
```

* Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|accessed_time

|string

a|Last access time of the file in date-time format.

|analytics

|link:#analytics[analytics]

a|Additional file system analytics information summarizing all descendents of a directory.

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a `xref:{relative_path}file-info-response(#model-file-info-response)`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the `file-info-response(#model-file-info-response)`. This avoids an excessive amount of duplicated information when a `get-storage-volumes-files-.html<<model-file_info_response,file_info_response>>`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the `<<model-file_info_response,file_info_response>>`. This avoids an excessive amount of duplicated information when a [GET

`/storage/volumes/{volume.uuid}/files/{path}]` call returns a large collection.

|bytes_used

|integer

a|The actual number of bytes used on disk by this file. If `byte_offset` and `length` parameters are specified, this will return the bytes used by the file within the given range.

|changed_time

|string

a|Last time data or attributes changed on the file in date-time format.

|constituent

|link:#constituent[constituent]

a|

|creation_time

|string

a|Creation time of the file in date-time format.

|fill_enabled

|boolean

a|Returns "true" if the space reservation is enabled. The field `overwrite_enabled` must also be set to the same value as this field.

|group_id

|integer

a|The integer ID of the group of the file owner.

|hard_links_count

|integer

a|The number of hard links to the file.

|inode_generation

|integer

a|Inode generation number.

|inode_number

|integer

a|The file inode number.

|is_empty

|boolean

a|Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set.

|is_junction

|boolean

a|Returns "true" if the directory is a junction.

|is_snapshot

|boolean

a|Returns "true" if the directory is a Snapshot copy.

|is_vm_aligned

|boolean

a|Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume.

|modified_time

|string

a|Last data modification time of the file in date-time format.

|name

|string

a|Name of the file.

|overwrite_enabled

|boolean

a|Returns "true" if the space reservation for overwrites is enabled. The field fill_enabled must also be set to the same value as this field.

|owner_id
|integer
a|The integer ID of the file owner.

|path
|string
a|Path of the file.

|qos_policy
|link:#qos_policy[qos_policy]
a|The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".

NOTE: Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

|size
|integer
a|The size of the file, in bytes.

|target
|string
a|The relative or absolute path contained in a symlink, in the form +++<some>+++ /+++<path>+++ .+++</path>++++++</some>+++

|type
|string
a|Type of the file.

|unique_bytes
|integer
a|Number of bytes uniquely held by this file. If byte_offset and length parameters are specified, this will return bytes uniquely held by the file within the given range.

```
|unix_permissions
```

```
|integer
```

a|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). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group.

```
|volume
```

```
|link:#volume[volume]
```

```
a|
```

```
|===
```

```
.Example request
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{
  "_links": {
    "metadata": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "accessed_time": "2019-06-12 15:00:16 +0000",
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": "2019-07",
        "oldest_label": "2019-07",

```

```
"percentages": [
  0.1,
  11.24,
  0.18,
  15.75,
  0.75,
  83.5,
  0
],
"values": [
  15925248,
  1735569408,
  27672576,
  2430595072,
  116105216,
  12889948160,
  0
]
},
"by_modified_time": {
  "bytes_used": {
    "labels": [
      "2019-07",
      "2019-06",
      "2019-05",
      "2019",
      "2018",
      "--2017",
      "unknown"
    ],
    "newest_label": "2019-07",
    "oldest_label": "2019-07",
    "percentages": [
      0.1,
      11.24,
      0.18,
      15.75,
      0.75,
      83.5,
      0
    ],
    "values": [
      15925248,
      1735569408,
      27672576,
```

```

        2430595072,
        116105216,
        12889948160,
        0
    ]
}
},
"bytes_used": 15436648448,
"file_count": 21134,
"subdir_count": 35
},
"bytes_used": 4096,
"changed_time": "2019-06-12 15:00:16 +0000",
"constituent": {
    "name": "fg__0001",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"creation_time": "2019-06-12 15:00:16 +0000",
"group_id": 30,
"hard_links_count": 1,
"inode_generation": 214753547,
"inode_number": 1695,
"is_empty": "",
"is_junction": "",
"is_snapshot": "",
"is_vm_aligned": "",
"modified_time": "2019-06-12 15:00:16 +0000",
"name": "test_file",
"owner_id": 54738,
"path": "d1/d2/d3",
"qos_policy": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    }
},
"name": "qos1",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": 200,
"target": "some_directory/some_other_directory/some_file",
"type": "file",
"unique_bytes": 4096,
"unix_permissions": 493,
"volume": {
    "_links": {

```

```
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
====

== Response
```

Status: 201, Created

```
==== Headers

[cols=3*,options=header]
|====
//header
|Name
|Description
|Type
//end header

//start row
|Location
|Useful for tracking the resource location
|string
//end row
//end table
|====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|====
| Error Code | Description

| 917505
| The SVM does not exist.

| 917525
```

```
| The volume in the symlink path does not exist in the SVM.

| 917698
| The volume in the symlink path is not mounted in the namespace.

| 6488064
| This command is not supported.

| 6488065
| The volume in the symlink path is invalid.

| 6488066
| Mounting the unjunctioned volume in the symlink path failed.

| 6488069
| Internal file error.

| 6488084
| Failed to create \{path} because the "unix_permissions" field was not
specified.

| 6488085
| Failed to create \{path} because the "type" field was not specified.

| 8257536
| This operation is not supported for the system volume specified in the
symlink path.

| 8257541
| Failed to compute the SVM identification from this content.

| 8257542
| This operation is not supported for the administrative SVM.

| 9437549
| This operation is not allowed on SVMs with Infinite Volume.

| 13172837
| This operation is not permitted because the SVM is locked for a migrate
operation.

| 111411203
| Failed to get the volume file system analytics report on the volume.

| 111411204
| Internal error. Failed to retrieve the volume file system analytics
report on the volume.
```



```
| 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.

| 111411215
| Internal error. Volume file system analytics report timed out for volume
volume.name in SVM svm.name.

| 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.

| 111412203
| Volume file system analytics is not enabled on the volume.
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```

|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|metadata
|link:#href[href]
a|
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#bytes_used]
[.api-collapsible-fifth-title]
bytes_used
```

Number of bytes used on-disk, broken down by date of last access.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|labels
|array[string]
a|Labels for this histogram
```

```
|newest_label
|string
a|Each label indicates the period of time the corresponding data is
associated with. A label can take one of the following forms:<ul>
  a partial date in an extended ISO8601 representation
  an interval between partial dates in an extended ISO8601 representation,
where "--" is used to separate the beginning and end of the interval
  the string literal "unknown"
```

For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00--2017-12-31T23:59:59. Similarly, the interval "2018-05--2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00--2018-07-31T23:59:59.

The following extensions to ISO8601 are used:

Quarters may be specified. The form `__yyyy_-Q__q__` is used to represent the `__q__`th quarter of the year `__yyyy_`. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00--2019-06-30T23:59:59.

Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval includes points in time through the end of the current week.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

|oldest_label

|string

a|Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:

- a partial date in an extended ISO8601 representation
- an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval
- the string literal "unknown"

For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00--2017-12-31T23:59:59. Similarly, the interval "2018-05--2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00--2018-07-31T23:59:59.

The following extensions to ISO8601 are used:

Quarters may be specified. The form `_yyyy_-Q__q__` is used to represent the `__q__`th quarter of the year `_yyyy_`. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00--2019-06-30T23:59:59.

Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval includes points in time through the end of the current week.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

```
|percentages
|array[number]
a|Percentages for this histogram
```

```
|values
|array[integer]
a|Values for this histogram
```

```
|===
```

```
[#by_accessed_time]
[.api-collapsible-fifth-title]
by_accessed_time
```

File system analytics information, broken down by date of last access.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|bytes_used
```

```
|link:#bytes_used[bytes_used]
```

```
a|Number of bytes used on-disk, broken down by date of last access.
```

```
|===
```

```
[#bytes_used]  
[.api-collapsible-fifth-title]  
bytes_used
```

Number of bytes used on-disk, broken down by date of last modification.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|labels
```

```
|array[string]
```

```
a|Labels for this histogram
```

```
|newest_label
```

```
|string
```

```
a|Each label indicates the period of time the corresponding data is  
associated with. A label can take one of the following forms:<ul>  
  a partial date in an extended ISO8601 representation  
  an interval between partial dates in an extended ISO8601 representation,  
  where "--" is used to separate the beginning and end of the interval  
  the string literal "unknown"
```

```
</ul>For partial dates and partial date intervals where components of a  
date are unspecified, the label allows for any valid normalized values the  
unspecified components might take. For example, the label "2017" allows  
for any time within the year 2017. Essentially, this is the fully  
specified interval 2017-01-01T00:00:00--2017-12-31T23:59:59. Similarly,  
the interval "2018-05--2018-07" allows for any time within the months of  
May, June, and July in 2018, corresponding to the fully specified interval  
2018-05-01T00:00:00--2018-07-31T23:59:59.
```

```
The following extensions to ISO8601 are used:<ul>
```

```
  Quarters may be specified. The form yyyy-Qq is used to represent  
  the qth quarter of the year yyyy. Q1 consists of the months January,  
  February, and March; Q2 consists of April, May, and June; Q3 consists of  
  July, August, and September; Q4 consists of October, November, and  
  December. For example, the label "2019-Q2" represents the second quarter
```

of the year 2019, which corresponds to the interval 2019-04-01T00:00:00--2019-06-30T23:59:59.

Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval includes points in time through the end of the current week.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

|oldest_label

|string

a|Each label indicates the period of time the corresponding data is associated with. A label can take one of the following forms:

- a partial date in an extended ISO8601 representation
- an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval
- the string literal "unknown"

For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00--2017-12-31T23:59:59. Similarly, the interval "2018-05--2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00--2018-07-31T23:59:59.

The following extensions to ISO8601 are used:

Quarters may be specified. The form `__yyyy_-Q__q__` is used to represent the `__q__`th quarter of the year `__yyyy_`. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00--2019-06-30T23:59:59.

Either the beginning or end of an interval may be omitted. When the beginning is omitted, the interval includes points in time arbitrarily far in the past. When the end is omitted, the interval includes points in time through the end of the current week.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

```
|percentages
|array[number]
a|Percentages for this histogram
```

```
|values
|array[integer]
a|Values for this histogram
```

```
|===
```

```
[#by_modified_time]
[.api-collapsible-fifth-title]
by_modified_time
```

File system analytics information, broken down by date of last modification.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|bytes_used
|link:#bytes_used[bytes_used]
a|Number of bytes used on-disk, broken down by date of last modification.
```

```
|===
```

```
[#analytics]
[.api-collapsible-fifth-title]
analytics
```

Additional file system analytics information summarizing all descendents of a directory.

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a `xref:{relative_path}file-`


```
info-response(#model-file-info-
response),analyticsobjectswillonlyincludepropertiesthatmayvarybetweeneleme
ntswithinthecollection.forexample,theanalyticsobjectswillnotcontainhistogr
amlabels,sincethesamehistogramlabelsareusedforallelementswithinthecollecti
on.theinvariantinformationisinsteadavailableviatheanalyticspropertyofthefi
le-info-response(#model-file-info-
response).thisavoidsanexcessiveamountofduplicatedinformationwhenaget-
storage-volumes-files-.html<<model-
file_info_response,file_info_response>>, analytics objects will only
include properties that may vary between elements within the collection.
For example, the analytics objects will not contain histogram labels,
since the same histogram labels are used for all elements within the
collection. The invariant information is instead available via the
analytics property of the <<model-file_info_response,file_info_response>>.
This avoids an excessive amount of duplicated information when a [GET
/storage/volumes/{volume.uuid}/files/{path}] call returns a large
collection.
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|by_accessed_time
```

```
|link:#by_accessed_time[by_accessed_time]
```

```
a|File system analytics information, broken down by date of last access.
```

```
|by_modified_time
```

```
|link:#by_modified_time[by_modified_time]
```

```
a|File system analytics information, broken down by date of last
modification.
```

```
|bytes_used
```

```
|integer
```

```
a|Number of bytes used on-disk
```

```
|file_count
```

```
|integer
```

```
a|Number of descendants
```

```
|incomplete_data
```

```

|boolean
a|Returns true if data collection is incomplete for this directory tree.

|subdir_count
|integer
a|Number of sub directories

|===

[#constituent]
[.api-collapsible-fifth-title]
constituent

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|FlexGroup volume constituent name.

|uuid
|string
a|FlexGroup volume constituent UUID.

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]

```

```
a|
```

```
|===
```

```
[#qos_policy]  
[.api-collapsible-fifth-title]  
qos_policy
```

The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both ``qos_policy.uuid`` and ``qos_policy.name`` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property ``qos_policy.name`` in a PATCH request to an empty string `""` or `"none"`.

NOTE: Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on `/storage/luns` to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]
```

```
a|
```

```
|name  
|string
```

a|The name of the QoS policy. To remove the file from a QoS policy, set this property to an empty string `""` or set it to `"none"` in a PATCH request.

```
|uuid  
|string
```

a|The unique identifier of the QoS policy. Valid in PATCH.

```
|===
```

```

[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume.

|uuid
|string
a|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

|===

[#file_info]
[.api-collapsible-fifth-title]
file_info

Information about a single file.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links

```

```
|link:#_links[_links]
```

```
a|
```

```
|accessed_time
```

```
|string
```

```
a|Last access time of the file in date-time format.
```

```
|analytics
```

```
|link:#analytics[analytics]
```

```
a|Additional file system analytics information summarizing all descendents of a directory.
```

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a `xref:{relative_path}file-info-response(#model-file-info-response)`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the `file-info-response(#model-file-info-response)`. This avoids an excessive amount of duplicated information when a `get-storage-volumes-files-.html<<model-file_info_response,file_info_response>>`, analytics objects will only include properties that may vary between elements within the collection. For example, the analytics objects will not contain histogram labels, since the same histogram labels are used for all elements within the collection. The invariant information is instead available via the analytics property of the `<<model-file_info_response,file_info_response>>`. This avoids an excessive amount of duplicated information when a `[GET /storage/volumes/{volume.uuid}/files/{path}]` call returns a large collection.

```
|bytes_used
```

```
|integer
```

```
a|The actual number of bytes used on disk by this file. If byte_offset and length parameters are specified, this will return the bytes used by the file within the given range.
```

```
|changed_time
```

```
|string
```

```
a|Last time data or attributes changed on the file in date-time format.
```

```
|constituent
|link:#constituent[constituent]
a|
```

```
|creation_time
|string
a|Creation time of the file in date-time format.
```

```
|fill_enabled
|boolean
a|Returns "true" if the space reservation is enabled. The field
overwrite_enabled must also be set to the same value as this field.
```

```
|group_id
|integer
a|The integer ID of the group of the file owner.
```

```
|hard_links_count
|integer
a|The number of hard links to the file.
```

```
|inode_generation
|integer
a|Inode generation number.
```

```
|inode_number
|integer
a|The file inode number.
```

```
|is_empty
|boolean
a|Specifies whether or not a directory is empty. A directory is considered
empty if it only contains entries for "." and "..". This element is
present if the file is a directory. In some special error cases, such as
when the volume goes offline or when the directory is moved while
retrieving this info, this field might not get set.
```

```
|is_junction
|boolean
```

a|Returns "true" if the directory is a junction.

|is_snapshot

|boolean

a|Returns "true" if the directory is a Snapshot copy.

|is_vm_aligned

|boolean

a|Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume.

|modified_time

|string

a|Last data modification time of the file in date-time format.

|name

|string

a|Name of the file.

|overwrite_enabled

|boolean

a|Returns "true" if the space reservation for overwrites is enabled. The field fill_enabled must also be set to the same value as this field.

|owner_id

|integer

a|The integer ID of the file owner.

|path

|string

a|Path of the file.

|qos_policy

|link:#qos_policy[qos_policy]

a|The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy.

To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".

NOTE: Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

|size
|integer
a|The size of the file, in bytes.

|target
|string
a|The relative or absolute path contained in a symlink, in the form
+++<some>+++ /+++<path>+++ .+++</path>++++++</some>+++

|type
|string
a|Type of the file.

|unique_bytes
|integer
a|Number of bytes uniquely held by this file. If byte_offset and length parameters are specified, this will return bytes uniquely held by the file within the given range.

|unix_permissions
|integer
a|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). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group.

|volume
|link:#volume[volume]
a|


```
|===
```

```
[#error_arguments]  
[.api-collapsible-fifth-title]  
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|code  
|string  
a|Argument code
```

```
|message  
|string  
a|Message argument
```

```
|===
```

```
[#error]  
[.api-collapsible-fifth-title]  
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|arguments  
|array[link:#error_arguments[error_arguments]]  
a|Message arguments
```

```
|code  
|string  
a|Error code
```

```
|message
```

```
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID7c1e16fab08e6059828926b28e8b7fc]]
= Retrieve historical performance metrics for a volume
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/volumes/{volume.uuid}/metrics`#
```

```
*Introduced In:* 9.7
```

```
Retrieves historical performance metrics for a volume.
```

```
== Parameters
```

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|latency.read
|integer
|query
|False
a|Filter by latency.read
```

```
|latency.other
```

```
|integer  
|query  
|False  
a|Filter by latency.other
```

```
|latency.write  
|integer  
|query  
|False  
a|Filter by latency.write
```

```
|latency.total  
|integer  
|query  
|False  
a|Filter by latency.total
```

```
|throughput.read  
|integer  
|query  
|False  
a|Filter by throughput.read
```

```
|throughput.other  
|integer  
|query  
|False  
a|Filter by throughput.other
```

```
|throughput.write  
|integer  
|query  
|False  
a|Filter by throughput.write
```

```
|throughput.total  
|integer  
|query  
|False  
a|Filter by throughput.total
```

```
|duration
|string
|query
|False
a|Filter by duration
```

```
|flexcache.status
|string
|query
|False
a|Filter by flexcache.status
```

* Introduced in: 9.8

```
|flexcache.timestamp
|string
|query
|False
a|Filter by flexcache.timestamp
```

* Introduced in: 9.8

```
|flexcache.bandwidth_savings
|integer
|query
|False
a|Filter by flexcache.bandwidth_savings
```

* Introduced in: 9.9

```
|flexcache.duration
|string
|query
|False
a|Filter by flexcache.duration
```

* Introduced in: 9.8

```
|flexcache.cache_miss_percent
|integer
|query
```

```
|False  
a|Filter by flexcache.cache_miss_percent
```

```
* Introduced in: 9.8
```

```
|iops.read  
|integer  
|query  
|False  
a|Filter by iops.read
```

```
|iops.other  
|integer  
|query  
|False  
a|Filter by iops.other
```

```
|iops.write  
|integer  
|query  
|False  
a|Filter by iops.write
```

```
|iops.total  
|integer  
|query  
|False  
a|Filter by iops.total
```

```
|cloud.iops.read  
|integer  
|query  
|False  
a|Filter by cloud.iops.read
```

```
|cloud.iops.other  
|integer  
|query  
|False  
a|Filter by cloud.iops.other
```

```
|cloud.iops.write
|integer
|query
|False
a|Filter by cloud.iops.write
```

```
|cloud.iops.total
|integer
|query
|False
a|Filter by cloud.iops.total
```

```
|cloud.latency.read
|integer
|query
|False
a|Filter by cloud.latency.read
```

```
|cloud.latency.other
|integer
|query
|False
a|Filter by cloud.latency.other
```

```
|cloud.latency.write
|integer
|query
|False
a|Filter by cloud.latency.write
```

```
|cloud.latency.total
|integer
|query
|False
a|Filter by cloud.latency.total
```

```
|cloud.status
|string
|query
|False
```

```
a|Filter by cloud.status
```

```
|cloud.timestamp
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by cloud.timestamp
```

```
|cloud.duration
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by cloud.duration
```

```
|timestamp
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by timestamp
```

```
|status
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by status
```

```
|volume.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Unique identifier of the volume.
```

```
|interval
```

```
|string
```

```
|query
```

```
|False
```

```
a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.  
The period for each time range is as follows:
```

```
* 1h: Metrics over the most recent hour sampled over 15 seconds.
```

```
* 1d: Metrics over the most recent day sampled over 5 minutes.
```

```

* 1w: Metrics over the most recent week sampled over 30 minutes.
* 1m: Metrics over the most recent month sampled over 2 hours.
* 1y: Metrics over the most recent year sampled over a day.
* Default value: 1
* enum: ["1h", "1d", "1w", "1m", "1y"]

|return_timeout
|integer
|query
|False
a|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

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|return_records
|boolean
|query
|False

```


a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|===

== Response

Status: 200, Ok

```
[cols=3*,options=header]
```

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|num_records

|integer

a|Number of records

|records

|array[link:#records[records]]

a|

|===

.Example response

[%collapsible%closed]

====

```
[source,json,subs=+macros]
```

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
}
```

```

"num_records": 1,
"records": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "latency": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "status": "ok",
  "throughput": {
    "read": 200,
    "total": 1000,
    "write": 100
  },
  "timestamp": "2017-01-25 11:20:13 +0000"
}
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]

```

```
|===  
|Name  
|Type  
|Description  
  
|next  
|link:#href[href]  
a|
```

```
|self  
|link:#href[href]  
a|
```

```
|===
```

```
[#_links]  
[.api-collapsible-fifth-title]  
_links
```

```
[cols=3*,options=header]
```

```
|===  
|Name  
|Type  
|Description
```

```
|self  
|link:#href[href]  
a|
```

```
|===
```

```
[#iops]  
[.api-collapsible-fifth-title]  
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===  
|Name  
|Type  
|Description
```

```
|other
```

```
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#records]  
[.api-collapsible-fifth-title]  
records
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|duration  
|string  
a|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  
|link:#iops[iops]  
a|The rate of I/O operations observed at the storage object.
```

```
|latency  
|link:#latency[latency]  
a|The round trip latency in microseconds observed at the storage object.
```

```
|status  
|string  
a|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
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```



```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
:leveloffset: -1
```

```
= Manage volume Snapshot copies
```

```
:leveloffset: +1
```

```
[[IDc3a49982a7bc006070ee9cdb03e546ce]]
```

```
= Storage volumes volume.uuid snapshots endpoint overview
```

```
== Overview
```

A Snapshot copy is the view of the filesystem as it exists at the time when the Snapshot copy is created.

In ONTAP, different types of Snapshot copies are supported, such as scheduled Snapshot copies, user requested Snapshot copies, SnapMirror Snapshot copies, and so on.

ONTAP Snapshot copy APIs allow you to create, modify, delete and retrieve Snapshot copies.

ONTAP Bulk Snapshot copy APIs allow you to create, modify, delete and retrieve Snapshot copies on multiple volumes in one request.

== Snapshot copy APIs

The following APIs are used to perform operations related to Snapshot copies.

```
&ndash; POST      /api/storage/volumes/{volume.uuid}/snapshots
&ndash; GET      /api/storage/volumes/{volume.uuid}/snapshots
&ndash; GET      /api/storage/volumes/{volume.uuid}/snapshots/{uid}
&ndash; PATCH    /api/storage/volumes/{volume.uuid}/snapshots/{uid}
&ndash; DELETE   /api/storage/volumes/{volume.uuid}/snapshots/{uid}
```

The following APIs are used to perform bulk operations related to Snapshot copies.

```
&ndash; POST      /api/storage/volumes/*/snapshots
&ndash; GET      /api/storage/volumes/*/snapshots
&ndash; PATCH    /api/storage/volumes/*/snapshots/{uid}
&ndash; DELETE   /api/storage/volumes/*/snapshots/{uid}
```

== Examples

=== Creating a Snapshot copy

The POST operation is used to create a Snapshot copy with the specified attributes.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots
```

```
# The call:
curl -X POST "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots" -H "accept:
application/hal+json" -d '{"name": "snapshot_copy", "comment": "Store this
copy." }'

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 13 Mar 2019 22:43:34 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/?name=snapshot_copy
Content-Length: 189
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "volume": {
        "name": "v2"
      },
      "svm": {
        "uuid": "8139f958-3c6e-11e9-a45f-005056bbc848",
        "name": "vs0"
      },
      "name": "snapshot_copy",
      "comment": "Store this copy."
    }
  ],
  "job": {
    "uuid": "6f68c85b-45e1-11e9-8fc7-005056bbc848",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/6f68c85b-45e1-11e9-8fc7-005056bbc848"
      }
    }
  }
}

# The Job:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 22:43:57 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
```

```
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 224
Content-Type: application/json
{
  "uuid": "6f68c85b-45e1-11e9-8fc7-005056bbc848",
  "description": "POST /api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/?name=snapshot_copy",
  "state": "success",
  "message": "success",
  "code": 0
}
-----
```

=== Retrieving Snapshot copy attributes

The GET operation is used to retrieve Snapshot copy attributes.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots
```

```
# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots" -H "accept:
application/hal+json"
```

```
# The response:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 21:14:06 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Type: application/json
Transfer-Encoding: chunked
{
  "records": [
    {
      "uuid": "402b6c73-73a0-4e89-a58a-75ee0ab3e8c0",
      "name": "hourly.2019-03-13_1305",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0"
        }
      }
    }
  ]
}
```

```

},
{
  "uuid": "f0dd497f-efe8-44b7-a4f4-bdd3890bc0c8",
  "name": "hourly.2019-03-13_1405",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/f0dd497f-efe8-44b7-a4f4-bdd3890bc0c8"
    }
  }
},
{
  "uuid": "02701900-51bd-46b8-9c77-47d9a9e2ce1d",
  "name": "hourly.2019-03-13_1522",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/02701900-51bd-46b8-9c77-47d9a9e2ce1d"
    }
  }
}
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots"
  }
}
}
}
-----

```

=== Creating bulk Snapshot copies

The POST operation is used to create a Snapshot copy with the same name on multiple volumes in one request.

This operation accepts a volume UUID or volume name and SVM, and a Snapshot copy name.

This operation only supports SnapMirror label attributes to be added to Snapshot copies during creation.

```

# The API:
/api/storage/volumes/*/snapshots

```

```
# The call:
curl -k -u admin -X POST "https://<mgmt-
ip>/api/storage/volumes/*/snapshots" -H 'accept: application/hal+json' -d
'{"records": [{"volume.uuid": "e8815adb-5209-11ec-b4ad-005056bbc3e8",
"name": "snapshot_copy"}, {"volume.uuid": "efda9101-5209-11ec-b4ad-
005056bbc3e8", "name": "snapshot_copy"}]}'

# The response:
HTTP/1.1 202 Accepted
Date: Tue, 14 Dec 2021 20:18:13 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Location: /api/storage/volumes/*/snapshots/
Content-Length: 209
Content-Type: application/json
{
  "num_records": 2,
  "job": {
    "uuid": "f7130fc0-5d1a-11ec-b78c-005056bbb467",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f7130fc0-5d1a-11ec-b78c-005056bbb467"
      }
    }
  }
}

# The Job:
HTTP/1.1 200 OK
Date: Tue, 14 Dec 2021 20:20:54 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: 258
Content-Type: application/json
Vary: Accept-Encoding
{
  "uuid": "f7130fc0-5d1a-11ec-b78c-005056bbb467",
  "description": "POST /api/storage/volumes/*/snapshots/",
```

```
"state": "success",
"message": "success",
"code": 0,
"start_time": "2021-12-14T15:18:13-05:00",
"end_time": "2021-12-14T15:18:13-05:00"
}
```

=== Retrieving Snapshot copy advanced attributes

A collection GET request is used to calculate the amount of Snapshot copy reclaimable space.

When the advanced privilege field 'reclaimable space' is requested, the API returns the amount of reclaimable space for the queried list of Snapshot copies.

The API:

```
/api/storage/volumes/{volume.uuid}/snapshots?fields=reclaimable_space
```

The call:

```
curl -X GET "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots?fields=reclaimable_space&name=hourly.2019-03-13_1305&#124;hourly.2019-03-13_1405&#124;hourly.2019-03-13_1522" -H "accept: application/hal+json"
```

The response:

```
HTTP/1.1 200 OK
```

```
Date: Wed, 13 Mar 2019 21:14:06 GMT
```

```
Server: libzapid-httpd
```

```
X-Content-Type-Options: nosniff
```

```
Cache-Control: no-cache,no-store,must-revalidate
```

```
Content-Type: application/json
```

```
Transfer-Encoding: chunked
```

```
{
"records": [
  {
    "uuid": "402b6c73-73a0-4e89-a58a-75ee0ab3e8c0",
    "name": "hourly.2019-03-13_1305",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0"
      }
    }
  }
]
```

```

},
{
  "uuid": "f0dd497f-efe8-44b7-a4f4-bdd3890bc0c8",
  "name": "hourly.2019-03-13_1405",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/f0dd497f-efe8-44b7-a4f4-bdd3890bc0c8"
    }
  }
},
{
  "uuid": "02701900-51bd-46b8-9c77-47d9a9e2ce1d",
  "name": "hourly.2019-03-13_1522",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/02701900-51bd-46b8-9c77-47d9a9e2ce1d"
    }
  }
}
],
"num_records": 3,
"reclaimable_space": 1567832,
"_links": {
  "self": {
    "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots"
  }
}
}
}
-----

```

=== Retrieving Snapshot copy advanced attributes

A collection GET request is used to calculate the delta between two Snapshot copies.

When the advanced privilege field 'delta' is requested, the API returns the delta between the queried Snapshot copies.

The API:

/api/storage/volumes/{volume.uuid}/snapshots?fields=delta

The call:


```

curl -X GET "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots?fields=delta&name=hourly.2
022-06-29_1105,hourly.2022-06-29_1205" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Connection: Keep-Alive
Content-Encoding: gzip
Content-Length: 378
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-
ancestors: 'self'
Content-Type: application/hal+json
Date: Wed, 29 Jun 2022 18:37:11 GMT
Keep-Alive: timeout=5, max=99
Server: libzapid-httpd
Vary: Accept-Encoding
X-Content-Type-Options: nosniff
{
"records": [
  {
    "uuid": "52a2247a-7735-4a92-bc3c-e51df1fe502f",
    "name": "hourly.2022-06-29_1105",
    "delta": {
      "size_consumed": 675840,
      "time_elapsed": "PT3H27M45S"
    }
  },
  {
    "uuid": "b399eb34-44fe-4689-9fb5-c8f72162dd77",
    "name": "hourly.2022-06-29_1205",
    "delta": {
      "size_consumed": 507904,
      "time_elapsed": "PT2H27M45S"
    }
  }
],
"num_records": 2,
"delta": {
  "size_consumed": 167936,
  "time_elapsed": "PT1H"
}
}
----

=== Retrieving the attributes of a specific Snapshot copy

```

The GET operation is used to retrieve the attributes of a specific Snapshot copy.

The API:

```
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}
```

The call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-  
acb6-005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0" -H  
"accept: application/hal+json"
```

The response:

```
HTTP/1.1 200 OK
```

```
Date: Wed, 13 Mar 2019 22:39:26 GMT
```

```
Server: libzapid-httpd
```

```
X-Content-Type-Options: nosniff
```

```
Cache-Control: no-cache,no-store,must-revalidate
```

```
Content-Length: 308
```

```
Content-Type: application/json
```

```
{  
  "volume": {  
    "uuid": "0353dc05-405f-11e9-acb6-005056bbc848",  
    "name": "v2",  
    "_links": {  
      "self": {  
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848"  
      }  
    }  
  },  
  "uuid": "402b6c73-73a0-4e89-a58a-75ee0ab3e8c0",  
  "svm": {  
    "uuid": "8139f958-3c6e-11e9-a45f-005056bbc848",  
    "name": "vs0",  
    "_links": {  
      "self": {  
        "href": "/api/svm/svms/8139f958-3c6e-11e9-a45f-005056bbc848"  
      }  
    }  
  },  
  "name": "hourly.2019-03-13_1305",  
  "create_time": "2019-03-13T13:05:00-04:00",  
  "size": 122880,  
  "_links": {
```

```
    "self": {
      "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0"
    }
  }
}
-----
```

=== Retrieving the advanced attributes of a specific Snapshot copy

The GET operation is used to retrieve the attributes of a specific Snapshot copy. Snapshot copy reclaimable space can be requested during a GET request.

When the advanced privilege field reclaimable space is requested, the API returns the amount of reclaimable space for the Snapshot copy.

The API:

```
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}?fields=reclaimable_space
```

The call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0?fields=*" -H "accept: application/hal+json"
```

The response:

```
HTTP/1.1 200 OK
```

```
Date: Wed, 13 Mar 2019 22:39:26 GMT
```

```
Server: libzapid-httpd
```

```
X-Content-Type-Options: nosniff
```

```
Cache-Control: no-cache,no-store,must-revalidate
```

```
Content-Length: 308
```

```
Content-Type: application/json
```

```
{
  "volume": {
    "uuid": "0353dc05-405f-11e9-acb6-005056bbc848",
    "name": "v2",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848"
      }
    }
  },
  "uuid": "402b6c73-73a0-4e89-a58a-75ee0ab3e8c0",
```

```
"svm": {
  "uuid": "8139f958-3c6e-11e9-a45f-005056bbc848",
  "name": "vs0",
  "_links": {
    "self": {
      "href": "/api/svm/svms/8139f958-3c6e-11e9-a45f-005056bbc848"
    }
  }
},
"name": "hourly.2019-03-13_1305",
"reclaimable_space": 167832,
"_links": {
  "self": {
    "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0"
  }
}
}
```

=== Retrieving Snapshot copy advanced attributes

A collection GET request is used to calculate the delta between two Snapshot copies.

When the advanced privilege field 'delta' is requested, the API returns the delta between the queried Snapshot copies.

The API:

```
/api/storage/volumes/{volume.uuid}/snapshots?fields=delta
```

The call:

```
curl -X GET "https://<mgmt-  
ip>/api/storage/volumes/{volume.uuid}/snapshots?fields=delta&name=hourly.2  
022-06-29_1105,hourly.2022-06-29_1205" -H "accept: application/hal+json"
```

The response:

```
HTTP/1.1 200 OK
```

```
Connection: Keep-Alive
```

```
Content-Encoding: gzip
```

```
Content-Length: 378
```

```
Content-Security-Policy: default-src 'self'; script-src 'self' 'unsafe-  
inline'; style-src 'self' 'unsafe-inline'; img-src 'self' data:; frame-  
ancestors: 'self'
```

```
Content-Type: application/hal+json
```

Date: Wed, 29 Jun 2022 18:37:11 GMT

Keep-Alive: timeout=5, max=99

Server: libzapid-httpd

Vary: Accept-Encoding

X-Content-Type-Options: nosniff

```
{
"records": [
  {
    "uuid": "52a2247a-7735-4a92-bc3c-e51df1fe502f",
    "name": "hourly.2022-06-29_1105",
    "delta": {
      "size_consumed": 675840,
      "time_elapsed": "PT3H27M45S"
    }
  },
  {
    "uuid": "b399eb34-44fe-4689-9fb5-c8f72162dd77",
    "name": "hourly.2022-06-29_1205",
    "delta": {
      "size_consumed": 507904,
      "time_elapsed": "PT2H27M45S"
    }
  }
],
"num_records": 2,
"delta": {
  "size_consumed": 167936,
  "time_elapsed": "PT1H"
}
}
-----
```

=== Retrieving bulk Snapshot copies

The bulk GET operation is used to retrieve Snapshot copy attributes across all volumes.

The API:

/api/storage/volumes/*/snapshots

The call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/*/snapshots" -H
"accept: application/hal+json"
```

```
# The response:
HTTP/1.1 200 OK
{
  "records": [
    {
      "volume": {
        "uuid": "966c285f-47f7-11ec-8407-005056bbc08f",
        "name": "v1"
      },
      "uuid": "3edba912-5507-4535-adce-e12fe5c0e31c",
      "name": "daily.2021-11-18_0010"
    },
    {
      "volume": {
        "uuid": "966c285f-47f7-11ec-8407-005056bbc08f",
        "name": "v1"
      },
      "uuid": "3ad61153-d5ef-495d-8e0e-5c3b8bbaf5e6",
      "name": "hourly.2021-11-18_0705"
    },
    {
      "volume": {
        "uuid": "99c974e3-47f7-11ec-8407-005056bbc08f",
        "name": "v2"
      },
      "uuid": "3dd0fa97-65d9-41ea-a99d-5ceb9d2f55c5",
      "name": "daily.2021-11-18_0010"
    },
    {
      "volume": {
        "uuid": "99c974e3-47f7-11ec-8407-005056bbc08f",
        "name": "v2"
      },
      "uuid": "6ca20a52-c342-4753-8865-3693fa9b7e23",
      "name": "hourly.2021-11-18_0705"
    },
  ],
  "num_records": 4
}
----
```

=== Updating a Snapshot copy

The PATCH operation is used to update the specific attributes of a Snapshot copy.

```
-----  
  
# The API:  
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}  
  
# The call:  
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-  
acb6-005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f" -d  
'{"name": "snapshot_copy_new" }' -H "accept: application/hal+json"  
  
# The response:  
HTTP/1.1 202 Accepted  
Date: Wed, 13 Mar 2019 22:50:44 GMT  
Server: libzapid-httpd  
X-Content-Type-Options: nosniff  
Cache-Control: no-cache,no-store,must-revalidate  
Content-Length: 189  
Content-Type: application/json  
{  
  "job": {  
    "uuid": "6f7c3a82-45e2-11e9-8fc7-005056bbc848",  
    "_links": {  
      "self": {  
        "href": "/api/cluster/jobs/6f7c3a82-45e2-11e9-8fc7-005056bbc848"  
      }  
    }  
  }  
}  
  
# The Job:  
HTTP/1.1 200 OK  
Date: Wed, 13 Mar 2019 22:54:16 GMT  
Server: libzapid-httpd  
X-Content-Type-Options: nosniff  
Cache-Control: no-cache,no-store,must-revalidate  
Content-Length: 242  
Content-Type: application/json  
{  
  "uuid": "6f7c3a82-45e2-11e9-8fc7-005056bbc848",  
  "description": "PATCH /api/storage/volumes/0353dc05-405f-11e9-acb6-  
005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f",  
  "state": "success",  
  "message": "success",  
  "code": 0  
}  
-----
```

=== Updating bulk Snapshot copies

The bulk PATCH operation is used to update the specific attributes of Snapshot copies across volumes in a single request.

The API:

```
/api/storage/volumes/*/snapshots
```

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/*/snapshots" -d
'{"records": [{"volume.uuid":"e8815adb-5209-11ec-b4ad-005056bbc3e8",
"svm.uuid":"d0e6def5-5209-11ec-b4ad-005056bbc3e8", "uuid":"f9b7714d-1166-
410a-b143-874f27969db6", "comment":"yay"}, {"volume.uuid":"efda9101-5209-
11ec-b4ad-005056bbc3e8", "svm.uuid":"d0e6def5-5209-11ec-b4ad-
005056bbc3e8", "uuid":"514c82a7-bff7-48e2-a13c-5337b09ed41e",
"comment":"yay"}]}' -H "accept: application/hal+json"
```

The response:

```
HTTP/1.1 202 Accepted
```

```
{
"job": {
  "uuid": "1e9a561f-520f-11ec-b4ad-005056bbc3e8",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/1e9a561f-520f-11ec-b4ad-005056bbc3e8"
    },
    "results": {
      "href": "/api/storage/volumes/*/snapshots?job_results_uuid=1e9a561f-
520f-11ec-b4ad-005056bbc3e8"
    }
  }
}
}
```

The Job:

```
curl -u admin:netappl! -k -X GET --header 'Content-Type: application/json'
--header 'Accept: application/json' 'https://<mgmt-
ip>/api/storage/volumes/*/snapshots?job_results_uuid=1e9a561f-520f-11ec-
b4ad-005056bbc3e8'
```

```
HTTP/1.1 200 OK
```

```
{
"records": [
{
```



```

    "volume": {
      "uuid": "e8815adb-5209-11ec-b4ad-005056bbc3e8",
      "name": "v1"
    },
    "uuid": "f9b7714d-1166-410a-b143-874f27969db6",
    "svm": {
      "uuid": "d0e6def5-5209-11ec-b4ad-005056bbc3e8"
    },
    "name": "s1",
    "comment": "yay"
  },
  {
    "volume": {
      "uuid": "efda9101-5209-11ec-b4ad-005056bbc3e8",
      "name": "v2"
    },
    "uuid": "514c82a7-bff7-48e2-a13c-5337b09ed41e",
    "svm": {
      "uuid": "d0e6def5-5209-11ec-b4ad-005056bbc3e8"
    },
    "name": "s1",
    "comment": "yay"
  }
],
"num_records": 2
}
-----

```

=== Deleting a Snapshot copy

The DELETE operation is used to delete a Snapshot copy.

The API:

/api/storage/volumes/{volume.uuid}/snapshots/{uuid}

The call:

```

curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-
acb6-005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f" -H
"accept: application/hal+json"

```

The response:

```

HTTP/1.1 202 Accepted
Date: Wed, 13 Mar 2019 22:57:51 GMT
Server: libzapid-httpd

```

```
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
  "job": {
    "uuid": "6da1dfdd-45e3-11e9-8fc7-005056bbc848",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/6da1dfdd-45e3-11e9-8fc7-005056bbc848"
      }
    }
  }
}
```

The Job:

```
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 23:02:46 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 243
Content-Type: application/json
{
  "uuid": "6da1dfdd-45e3-11e9-8fc7-005056bbc848",
  "description": "DELETE /api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f",
  "state": "success",
  "message": "success",
  "code": 0
}
----
```

=== Deleting bulk Snapshot copies

The bulk DELETE operation is used to delete a Snapshot copies across volumes in a single request.

The API:

```
/api/storage/volumes/*/snapshots
```

The call:

```
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/*/snapshots" -d
'{"records": [{"volume.uuid": "e8815adb-5209-11ec-b4ad-005056bbc3e8",
```

```

"uuid":"f9b7714d-1166-410a-b143-874f27969db6"}, {"volume.uuid":"efda9101-
5209-11ec-b4ad-005056bbc3e8", "uuid":"1d55c97a-25f3-4366-bfa8-
9ea75c255469"}]}' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
{
  "job": {
    "uuid": "fe114ed7-520f-11ec-b4ad-005056bbc3e8",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/fe114ed7-520f-11ec-b4ad-005056bbc3e8"
      },
      "results": {
        "href": "/api/storage/volumes/*/snapshots?job_results_uuid=fe114ed7-
520f-11ec-b4ad-005056bbc3e8"
      }
    }
  }
}

# The Job:
HTTP/1.1 200 OK
curl -u admin:netappl! -k -X GET --header 'Content-Type: application/json'
--header 'Accept: application/json' 'https://<mgmt-
ip>/api/storage/volumes/*/snapshots?job_results_uuid=fe114ed7-520f-11ec-
b4ad-005056bbc3e8'
{
  "records": [
  ],
  "num_records": 0
}
-----

[[IDa30d9b9b677b8fa691fa5c1be9804605]]
= Retrieve volume Snapshot copies

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/volumes/{volume.uuid}/snapshots`#

*Introduced In:* 9.6

Retrieves a collection of volume Snapshot copies.

```

== Expensive properties

There is an added computational cost to retrieving the amount of reclaimable space for Snapshot copies, as the calculation is done on demand based on the list of Snapshot copies provided.

- * `reclaimable_space`
- * `delta`

== Related ONTAP commands

- * `snapshot show`
- * `snapshot compute-reclaimable`
- * `snapshot show-delta`

== Learn more

*
xref:{relative_path}storage_volumes_volume.uuid_snapshots_endpoint_overview.html[DOC /storage/volumes/{volume.uuid}/snapshots]

== Parameters

[cols=5*,options=header]

|===

|Name
|Type
|In
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume

|create_time
|string
|query
|False
a|Filter by create_time

```
|state
|string
|query
|False
a|Filter by state
```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|snapmirror_label
|string
|query
|False
a|Filter by snapmirror_label
```

* Introduced in: 9.8

```
|delta.size_consumed
|integer
|query
|False
a|Filter by delta.size_consumed
```

* Introduced in: 9.12

```
|delta.time_elapsed
|string
|query
|False
a|Filter by delta.time_elapsed
```

* Introduced in: 9.12

```
|owners  
|string  
|query  
|False  
a|Filter by owners
```

* Introduced in: 9.7

```
|logical_size  
|integer  
|query  
|False  
a|Filter by logical_size
```

* Introduced in: 9.12

```
|comment  
|string  
|query  
|False  
a|Filter by comment
```

```
|volume.name  
|string  
|query  
|False  
a|Filter by volume.name
```

```
|reclaimable_space  
|integer  
|query  
|False  
a|Filter by reclaimable_space
```

* Introduced in: 9.10

```
|expiry_time  
|string  
|query  
|False
```

a|Filter by expiry_time

|size

|integer

|query

|False

a|Filter by size

* Introduced in: 9.9

|version_uuid

|string

|query

|False

a|Filter by version_uuid

* Introduced in: 9.11

|snaplock_expiry_time

|string

|query

|False

a|Filter by snaplock_expiry_time

|provenance_volume.uuid

|string

|query

|False

a|Filter by provenance_volume.uuid

* Introduced in: 9.11

|svm.uuid

|string

|query

|False

a|Filter by svm.uuid

|svm.name

|string

|query

```
|False
a|Filter by svm.name

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|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
a|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
a|Order results by specified fields and optional [asc|desc] direction.
```


Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

```
[cols=3*,options=header]
```

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|delta

|link:#snapshot_delta[snapshot_delta]

a|Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a Snapshot copy and the Active File System, only the Snapshot copy name needs to be mentioned.

|num_records

|integer

a|Number of records

|reclaimable_space

|integer

a|Space reclaimed when the Snapshot copy is deleted, in bytes.

|records

|array[link:#snapshot[snapshot]]

a|

|===

.Example response

[%collapsible%closed]

====

```
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "delta": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "size_consumed": 0,
    "time_elapsed": "string"
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "comment": "string",
    "create_time": "2019-02-04 19:00:00 +0000",
    "delta": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "size_consumed": 0,
      "time_elapsed": "string"
    },
    "expiry_time": "2019-02-04 19:00:00 +0000",
    "logical_size": 1228800,
    "name": "this_snapshot",
    "owners": {
    },
    "provenance_volume": {
      "uuid": "4cd8a442-86d1-11e0-ae1c-125648563413"
    }
  },
}
```

```

"size": "122880",
"snaplock_expiry_time": "2019-02-04 19:00:00 +0000",
"state": "valid",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"version_uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
}
====

== Error

```

Status: Default

ONTAP Error Response Code

```

|===
| Error Code | Description
| 918235
| The specified volume is invalid.
|===

```

[cols=3*,options=header]

```

|===
|Name
|Type

```

```

|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

```

```

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#snapshot_delta]
[.api-collapsible-fifth-title]
snapshot_delta

```

Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated

format using the "name" parameter. To determine the space consumed between a Snapshot copy and the Active File System, only the Snapshot copy name needs to be mentioned.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|size_consumed
```

```
|integer
```

```
a|Indicates the space that has changed between two specified WAFL file systems, in bytes.
```

```
|time_elapsed
```

```
|string
```

```
a|Time elapsed between two specified WAFL file systems.
```

```
|===
```

```
[#provenance_volume]
```

```
[.api-collapsible-fifth-title]
```

```
provenance_volume
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|uuid
```

```
|string
```

```
a|UUID for the volume that is used to identify the source volume in a mirroring relationship. When the mirroring relationship is broken, a volume's Instance UUID and Provenance UUID are made identical. An unmirrored volume's Provenance UUID is the same as its Instance UUID. This field is valid for flexible volumes only.
```

```

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string

```

a|The name of the volume.

|uuid

|string

a|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

|===

[#snapshot]

[.api-collapsible-fifth-title]

snapshot

The Snapshot copy object represents a point in time Snapshot copy of a volume.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|comment

|string

a|A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.

|create_time

|string

a|Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.

|delta
|link:#snapshot_delta[snapshot_delta]
a|Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a Snapshot copy and the Active File System, only the Snapshot copy name needs to be mentioned.

|expiry_time
|string
a|The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.

|logical_size
|integer
a|Size of the logical used file system at the time the Snapshot copy is captured.

|name
|string
a|Snapshot copy. Valid in POST or PATCH.

|owners
|array[string]
a|

|provenance_volume
|link:#provenance_volume[provenance_volume]
a|

|reclaimable_space
|integer
a|Space reclaimed when the Snapshot copy is deleted, in bytes.

|size
|integer
a|Size of the active file system at the time the Snapshot copy is captured. The actual size of the Snapshot copy also includes those blocks trapped by other Snapshot copies. On a Snapshot copy deletion, the "size" amount of blocks is the maximum number of blocks available. On a Snapshot copy restore, the "afs-used size" value will match the Snapshot copy

"size" value.

|snaplock_expiry_time

|string

a|SnapLock expiry time for the Snapshot copy, if the Snapshot copy is taken on a SnapLock volume. A Snapshot copy is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during Snapshot copy POST and Snapshot copy PATCH on Snapshot copy locking enabled volumes.

|snapmirror_label

|string

a|Label for SnapMirror operations

|state

|string

a|State of the Snapshot copy. There are cases where some Snapshot copies are not complete. In the "partial" state, the Snapshot copy is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial Snapshot copies cannot be used for a Snapshot copy restore operation. A Snapshot copy is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a Snapshot copy is valid.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|The UUID of the Snapshot copy in the volume that uniquely identifies the Snapshot copy in that volume.

|version_uuid

|string

a|The 128 bit identifier that uniquely identifies a snapshot and its logical data layout.

|volume

|link:#volume[volume]

a|

```
|===
```

```
[#error_arguments]  
[.api-collapsible-fifth-title]  
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|code  
|string  
a|Argument code
```

```
|message  
|string  
a|Message argument
```

```
|===
```

```
[#error]  
[.api-collapsible-fifth-title]  
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|arguments  
|array[link:#error_arguments[error_arguments]]  
a|Message arguments
```

```
|code  
|string  
a|Error code
```

```

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[IDf76babc927817a140a2be6256ee95a0b]]
= Create a volume Snapshot copy

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/volumes/{volume.uuid}/snapshots`#

*Introduced In:* 9.6

Creates a volume Snapshot copy.

== Required properties

* `name` - Name of the Snapshot copy to be created.

== Recommended optional properties

* `comment` - Comment associated with the Snapshot copy.
* `expiry_time` - Snapshot copies with an expiry time set are not allowed
to be deleted until the retention time is reached.
* `snapmirror_label` - Label for SnapMirror operations.
* `snaplock_expiry_time` - Expiry time for Snapshot copy locking enabled
volumes.

== Related ONTAP commands

* `snapshot create`

== Learn more

```

```

*
xref:{relative_path}storage_volumes_volume.uuid_snapshots_endpoint_overvie
w.html[DOC /storage/volumes/{volume.uuid}/snapshots]

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume UUID

|return_timeout
|integer
|query
|False
a|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
a|The default is false. If set to true, the records are returned.

* Default value:

```

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|comment
```

```
|string
```

```
a|A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.
```

```
|create_time
```

```
|string
```

```
a|Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.
```

```
|delta
```

```
|link:#snapshot_delta[snapshot_delta]
```

```
a|Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a Snapshot copy and the Active File System, only the Snapshot copy name needs to be mentioned.
```

```
|expiry_time
```

```
|string
```

```
a|The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.
```

```
|logical_size
```

```
|integer
```

a|Size of the logical used file system at the time the Snapshot copy is captured.

|name

|string

a|Snapshot copy. Valid in POST or PATCH.

|owners

|array[string]

a|

|provenance_volume

|link:#provenance_volume[provenance_volume]

a|

|reclaimable_space

|integer

a|Space reclaimed when the Snapshot copy is deleted, in bytes.

|size

|integer

a|Size of the active file system at the time the Snapshot copy is captured. The actual size of the Snapshot copy also includes those blocks trapped by other Snapshot copies. On a Snapshot copy deletion, the "size" amount of blocks is the maximum number of blocks available. On a Snapshot copy restore, the "afs-used size" value will match the Snapshot copy "size" value.

|snaplock_expiry_time

|string

a|SnapLock expiry time for the Snapshot copy, if the Snapshot copy is taken on a SnapLock volume. A Snapshot copy is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during Snapshot copy POST and Snapshot copy PATCH on Snapshot copy locking enabled volumes.

|snapmirror_label

|string

a|Label for SnapMirror operations

|state

```
|string
a|State of the Snapshot copy. There are cases where some Snapshot copies
are not complete. In the "partial" state, the Snapshot copy is consistent
but exists only on the subset of the constituents that existed prior to
the FlexGroup's expansion. Partial Snapshot copies cannot be used for a
Snapshot copy restore operation. A Snapshot copy is in an "invalid" state
when it is present in some FlexGroup constituents but not in others. At
all other times, a Snapshot copy is valid.
```

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|The UUID of the Snapshot copy in the volume that uniquely identifies the
Snapshot copy in that volume.
```

```
|version_uuid
|string
a|The 128 bit identifier that uniquely identifies a snapshot and its
logical data layout.
```

```
|volume
|link:#volume[volume]
a|
```

```
|===
```

```
.Example request
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "create_time": "2019-02-04 19:00:00 +0000",
  "delta": {
    "_links": {
```



```
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "size_consumed": 0,
  "time_elapsed": "string"
},
"expiry_time": "2019-02-04 19:00:00 +0000",
"logical_size": 1228800,
"name": "this_snapshot",
"owners": {
},
"provenance_volume": {
  "uuid": "4cd8a442-86d1-11e0-ae1c-125648563413"
},
"size": "122880",
"snaplock_expiry_time": "2019-02-04 19:00:00 +0000",
"state": "valid",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"version_uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volumel1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
====

== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|job
```

```
|link:#job_link[job_link]
```

```
a|
```

```
|===
```

```
.Example response
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{
```

```
  "job": {
```

```
    "_links": {
```

```
      "self": {
```

```
        "href": "/api/resourcelink"
```

```
      }
```

```
    },
```

```
    "uuid": "string"
```

```
  }
```

```
}
```

```
====
```

```
=== Headers
```

```
[cols=3*,options=header]
```

```
|===
```

```
//header
```

```
|Name
```

```
|Description
```

```
|Type
```

```
//end header
```

```
//start row
```

```
|Location
```

```
|Useful for tracking the resource location
```

```
|string
```

```
//end row
```

```
//end table
```

```
|===
```

Status: Default

ONTAP Error Response Code

|===

| Error Code | Description

| 524479

| The specified volume is not online or does not have enough space to create a Snapshot copy.

| 2621462

| The specified SVM name does not exist.

| 1638433

| A Snapshot copy with the specified name already exists.

| 1638461

| Snapshot copies can only be created on read/write (RW) volumes.

| 1638477

| User-created Snapshot copy names cannot begin with the specified prefix.

| 1638518

| The specified Snapshot copy name is invalid.

| 1638532

| Failed to create the Snapshot copy on the specified volume because a revert operation is in progress.

| 1638537

| Cannot determine the status of the Snapshot copy create operation for the specified volume.

| 1638616

| Bulk Snapshot copy create is not supported with multiple Snapshot copy names.

| 1638617

| Bulk Snapshot copy create is not supported with volume names in a mixed-version cluster.

| 1638618

| The property cannot be specified for Snapshot copy create.

```

|===

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===

```

```
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#snapshot_delta]
[.api-collapsible-fifth-title]
snapshot_delta
```

Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a Snapshot copy and the Active File System, only the Snapshot copy name needs to be mentioned.

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
```

a|

|size_consumed

|integer

a|Indicates the space that has changed between two specified WAFL file systems, in bytes.

|time_elapsed

|string

a|Time elapsed between two specified WAFL file systems.

|===

[#provenance_volume]

[.api-collapsible-fifth-title]

provenance_volume

[cols=3*,options=header]

|===

|Name

|Type

|Description

|uuid

|string

a|UUID for the volume that is used to identify the source volume in a mirroring relationship. When the mirroring relationship is broken, a volume's Instance UUID and Provenance UUID are made identical. An unmirrored volume's Provenance UUID is the same as its Instance UUID. This field is valid for flexible volumes only.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#volume]
[.api-collapsible-fifth-title]
volume
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the volume.
```

```
|uuid
|string
a|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
```

```
|===
```

```
[#snapshot]
```

```
[.api-collapsible-fifth-title]
```

```
snapshot
```

The Snapshot copy object represents a point in time Snapshot copy of a volume.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
 |_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|comment
```

```
|string
```

a|A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.

```
|create_time
```

```
|string
```

a|Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.

```
|delta
```

```
|link:#snapshot_delta[snapshot_delta]
```

a|Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a Snapshot copy and the Active File System, only the Snapshot copy name needs to be mentioned.

```
|expiry_time
```

```
|string
```

a|The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to

be deleted until the retention time is reached.

|logical_size

|integer

a|Size of the logical used file system at the time the Snapshot copy is captured.

|name

|string

a|Snapshot copy. Valid in POST or PATCH.

|owners

|array[string]

a|

|provenance_volume

|link:#provenance_volume[provenance_volume]

a|

|reclaimable_space

|integer

a|Space reclaimed when the Snapshot copy is deleted, in bytes.

|size

|integer

a|Size of the active file system at the time the Snapshot copy is captured. The actual size of the Snapshot copy also includes those blocks trapped by other Snapshot copies. On a Snapshot copy deletion, the "size" amount of blocks is the maximum number of blocks available. On a Snapshot copy restore, the "afs-used size" value will match the Snapshot copy "size" value.

|snaplock_expiry_time

|string

a|SnapLock expiry time for the Snapshot copy, if the Snapshot copy is taken on a SnapLock volume. A Snapshot copy is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during Snapshot copy POST and Snapshot copy PATCH on Snapshot copy locking enabled volumes.

|snapmirror_label

```
|string
a|Label for SnapMirror operations
```

```
|state
|string
a|State of the Snapshot copy. There are cases where some Snapshot copies are not complete. In the "partial" state, the Snapshot copy is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial Snapshot copies cannot be used for a Snapshot copy restore operation. A Snapshot copy is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a Snapshot copy is valid.
```

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|The UUID of the Snapshot copy in the volume that uniquely identifies the Snapshot copy in that volume.
```

```
|version_uuid
|string
a|The 128 bit identifier that uniquely identifies a snapshot and its logical data layout.
```

```
|volume
|link:#volume[volume]
a|
```

```
|===
```

```
[#job_link]
[.api-collapsible-fifth-title]
job_link
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDe50e49454a8a2c22420754f046623f7a]]
= Delete a volume Snapshot copy

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/storage/volumes/{volume.uuid}/snapshots/{uuid}`#

*Introduced In:* 9.6

Deletes a Volume Snapshot copy.

== Related ONTAP commands

* `snapshot delete`

== Learn more

*
xref:{relative_path}storage_volumes_volume.uuid_snapshots_endpoint_overvie
w.html[DOC /storage/volumes/{volume.uuid}/snapshots]

```

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|volume.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Volume UUID
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Snapshot copy UUID
```

```
|return_timeout
```

```
|integer
```

```
|query
```

```
|False
```

```
a|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: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Error
```

Status: Default

```
ONTAP Error Response Code

|===
| Error Code | Description

| 2
| An invalid value was entered for one of the fields.

| 1638521
| Snapshot copies can only be deleted on read/write (RW) volumes.
```

```
| 1638538
| Cannot determine the status of the Snapshot copy delete operation for
the specified volume.
```

```
| 1638543
| Failed to delete Snapshot copy because it has an owner.
```

```
| 1638555
| The specified Snapshot copy has not expired or is locked.
```

```
| 1638600
| The Snapshot copy does not exist.
```

```
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```
.Example error
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

```
====
```

```

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===

```



```

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name

```

```
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[IDc0bc39894c9a9ffd602e3547f34ad2]]
= Retrieve volume Snapshot copy details
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/volumes/{volume.uuid}/snapshots/{uuid}`#
```

```
*Introduced In:* 9.6
```

Retrieves details of a specific volume Snapshot copy.

```
== Related ONTAP commands
```

```
* `snapshot show`
```

```
== Learn more
```

```

*
xref:{relative_path}storage_volumes_volume.uuid_snapshots_endpoint_overvie
w.html[DOC /storage/volumes/{volume.uuid}/snapshots]

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume UUID

|uuid
|string
|path
|True
a|Snapshot copy UUID

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type

```

|Description

|_links

|link:#_links[_links]

a|

|comment

|string

a|A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.

|create_time

|string

a|Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.

|delta

|link:#snapshot_delta[snapshot_delta]

a|Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a Snapshot copy and the Active File System, only the Snapshot copy name needs to be mentioned.

|expiry_time

|string

a|The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.

|logical_size

|integer

a|Size of the logical used file system at the time the Snapshot copy is captured.

|name

|string

a|Snapshot copy. Valid in POST or PATCH.

|owners

|array[string]

a|

|provenance_volume
|link:#provenance_volume[provenance_volume]

a|

|reclaimable_space
|integer
a|Space reclaimed when the Snapshot copy is deleted, in bytes.

|size
|integer
a|Size of the active file system at the time the Snapshot copy is captured. The actual size of the Snapshot copy also includes those blocks trapped by other Snapshot copies. On a Snapshot copy deletion, the "size" amount of blocks is the maximum number of blocks available. On a Snapshot copy restore, the "afs-used size" value will match the Snapshot copy "size" value.

|snaplock_expiry_time
|string
a|SnapLock expiry time for the Snapshot copy, if the Snapshot copy is taken on a SnapLock volume. A Snapshot copy is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during Snapshot copy POST and Snapshot copy PATCH on Snapshot copy locking enabled volumes.

|snapmirror_label
|string
a|Label for SnapMirror operations

|state
|string
a|State of the Snapshot copy. There are cases where some Snapshot copies are not complete. In the "partial" state, the Snapshot copy is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial Snapshot copies cannot be used for a Snapshot copy restore operation. A Snapshot copy is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a Snapshot copy is valid.

|svm

```

|link:#svm[svm]
a|

|uuid
|string
a|The UUID of the Snapshot copy in the volume that uniquely identifies the
Snapshot copy in that volume.

|version_uuid
|string
a|The 128 bit identifier that uniquely identifies a snapshot and its
logical data layout.

|volume
|link:#volume[volume]
a|

|===

```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```

{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "create_time": "2019-02-04 19:00:00 +0000",
  "delta": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "size_consumed": 0,
    "time_elapsed": "string"
  },
  "expiry_time": "2019-02-04 19:00:00 +0000",
  "logical_size": 1228800,
  "name": "this_snapshot",

```

```

"owners": {
},
"provenance_volume": {
  "uuid": "4cd8a442-86d1-11e0-ae1c-125648563413"
},
"size": "122880",
"snaplock_expiry_time": "2019-02-04 19:00:00 +0000",
"state": "valid",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"version_uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volumel1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
====

== Error

```

Status: Default

ONTAP Error Response Code

```

|====
| Error Code | Description
|
| 2
| An invalid value was entered for one of the fields.
|
| 262197
| An invalid field was specified in the request.

```

```
| 1638473
| Snapshot copy tag not found.

| 1638503
| The Snapshot copy does not exist on the specified volume.

| 1638600
| The Snapshot copy does not exist.

| 1638615
| Bulk operations for Snapshot copies are not supported on multiple SVMs.
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```
.Example error
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

```
====
```

```
== Definitions
```



```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#snapshot_delta]
[.api-collapsible-fifth-title]
snapshot_delta

```

Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a Snapshot copy and the Active File System, only the Snapshot copy name

needs to be mentioned.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|size_consumed
```

```
|integer
```

```
a|Indicates the space that has changed between two specified WAFL file systems, in bytes.
```

```
|time_elapsed
```

```
|string
```

```
a|Time elapsed between two specified WAFL file systems.
```

```
|===
```

```
[#provenance_volume]
```

```
[.api-collapsible-fifth-title]
```

```
provenance_volume
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|uuid
```

```
|string
```

```
a|UUID for the volume that is used to identify the source volume in a mirroring relationship. When the mirroring relationship is broken, a volume's Instance UUID and Provenance UUID are made identical. An unmirrored volume's Provenance UUID is the same as its Instance UUID. This field is valid for flexible volumes only.
```

```
|===
```

```

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume.

```

```
|uuid
|string
a|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
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[IDeda56e082521eb1bd9bf7ce91f50e2a7]]
= Update a volume Snapshot copy

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/volumes/{volume.uuid}/snapshots/{uuid}`#

*Introduced In:* 9.6

Updates a Volume Snapshot copy.

== Related ONTAP commands

* `snapshot modify`
* `snapshot rename`

== Learn more

*

```

```
xref:{relative_path}storage_volumes_volume.uuid_snapshots_endpoint_overview.html[DOC /storage/volumes/{volume.uuid}/snapshots]
```

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|volume.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Volume UUID
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Snapshot copy UUID
```

```
|return_timeout
```

```
|integer
```

```
|query
```

```
|False
```

```
a|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

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
 |_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|comment
```

```
|string
```

a|A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.

```
|create_time
```

```
|string
```

a|Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.

```
|delta
```

```
|link:#snapshot_delta[snapshot_delta]
```

a|Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a Snapshot copy and the Active File System, only the Snapshot copy name needs to be mentioned.

```
|expiry_time
```

```
|string
```

a|The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.

```
|logical_size
```

```
|integer
```

a|Size of the logical used file system at the time the Snapshot copy is captured.

```

|name
|string
a|Snapshot copy. Valid in POST or PATCH.

|owners
|array[string]
a|

|provenance_volume
|link:#provenance_volume[provenance_volume]
a|

|reclaimable_space
|integer
a|Space reclaimed when the Snapshot copy is deleted, in bytes.

|size
|integer
a|Size of the active file system at the time the Snapshot copy is
captured. The actual size of the Snapshot copy also includes those blocks
trapped by other Snapshot copies. On a Snapshot copy deletion, the "size"
amount of blocks is the maximum number of blocks available. On a Snapshot
copy restore, the "afs-used size" value will match the Snapshot copy
"size" value.

|snaplock_expiry_time
|string
a|SnapLock expiry time for the Snapshot copy, if the Snapshot copy is
taken on a SnapLock volume. A Snapshot copy is not allowed to be deleted
or renamed until the SnapLock ComplianceClock time goes beyond this
retention time. This option can be set during Snapshot copy POST and
Snapshot copy PATCH on Snapshot copy locking enabled volumes.

|snapmirror_label
|string
a|Label for SnapMirror operations

|state
|string
a|State of the Snapshot copy. There are cases where some Snapshot copies
are not complete. In the "partial" state, the Snapshot copy is consistent

```


but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial Snapshot copies cannot be used for a Snapshot copy restore operation. A Snapshot copy is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a Snapshot copy is valid.

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|The UUID of the Snapshot copy in the volume that uniquely identifies the
Snapshot copy in that volume.
```

```
|version_uuid
|string
a|The 128 bit identifier that uniquely identifies a snapshot and its
logical data layout.
```

```
|volume
|link:#volume[volume]
a|
```

```
|===
```

.Example request

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "create_time": "2019-02-04 19:00:00 +0000",
  "delta": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  }
}
```

```
    },
    "size_consumed": 0,
    "time_elapsed": "string"
  },
  "expiry_time": "2019-02-04 19:00:00 +0000",
  "logical_size": 1228800,
  "name": "this_snapshot",
  "owners": {
  },
  "provenance_volume": {
    "uuid": "4cd8a442-86d1-11e0-ae1c-125648563413"
  },
  "size": "122880",
  "snaplock_expiry_time": "2019-02-04 19:00:00 +0000",
  "state": "valid",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "version_uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volumel1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Error

```

Status: Default

ONTAP Error Response Code

```

|===
| Error Code | Description

| 524508
| The Snapshot copy was not renamed because the name entered is not valid.

| 542797
| The specified file or Snapshot copy does not exist.

| 1638455

```

```
| Failed to set comment for Snapshot copy.

| 1638476
| You cannot rename a Snapshot copy created for use as a reference
Snapshot copy by other jobs.

| 1638477
| User-created Snapshot copy names cannot begin with the specified prefix.

| 1638518
| The specified Snapshot copy name is invalid.

| 1638522
| Snapshot copies can only be renamed on read/write (RW) volumes.

| 1638523
| Failed to set the specified SnapMirror label for the Snapshot copy.

| 1638524
| Adding SnapMirror labels is not allowed in a mixed version cluster.

| 1638539
| Cannot determine the status of the Snapshot copy rename operation for
the specified volume.

| 1638554
| Failed to set expiry time for the Snapshot copy.

| 1638600
| The Snapshot copy does not exist.
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]

```

```
|===  
|Name  
|Type  
|Description  
  
|self  
|link:#href[href]  
a|
```

```
|===
```

```
[#snapshot_delta]  
[.api-collapsible-fifth-title]  
snapshot_delta
```

Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a Snapshot copy and the Active File System, only the Snapshot copy name needs to be mentioned.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|size_consumed  
|integer
```

a|Indicates the space that has changed between two specified WAFL file systems, in bytes.

```
|time_elapsed  
|string
```

a|Time elapsed between two specified WAFL file systems.

```
|===
```

```

[#provenance_volume]
[.api-collapsible-fifth-title]
provenance_volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|UUID for the volume that is used to identify the source volume in a
mirroring relationship. When the mirroring relationship is broken, a
volume's Instance UUID and Provenance UUID are made identical. An
unmirrored volume's Provenance UUID is the same as its Instance UUID. This
field is valid for flexible volumes only.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

```

```
[#volume]
[.api-collapsible-fifth-title]
volume
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the volume.
```

```
|uuid
```

```
|string
```

```
a|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
```

```
|===
```

```
[#snapshot]
```

```
[.api-collapsible-fifth-title]
```

```
snapshot
```

The Snapshot copy object represents a point in time Snapshot copy of a volume.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```


|Description

|_links

|link:#_links[_links]

a|

|comment

|string

a|A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.

|create_time

|string

a|Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.

|delta

|link:#snapshot_delta[snapshot_delta]

a|Reports the amount of space consumed between two WAFL file systems, in bytes. The two WAFL file systems should be specified in a comma-separated format using the "name" parameter. To determine the space consumed between a Snapshot copy and the Active File System, only the Snapshot copy name needs to be mentioned.

|expiry_time

|string

a|The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.

|logical_size

|integer

a|Size of the logical used file system at the time the Snapshot copy is captured.

|name

|string

a|Snapshot copy. Valid in POST or PATCH.

|owners

|array[string]

a|

|provenance_volume
|link:#provenance_volume[provenance_volume]

a|

|reclaimable_space
|integer

a|Space reclaimed when the Snapshot copy is deleted, in bytes.

|size
|integer

a|Size of the active file system at the time the Snapshot copy is captured. The actual size of the Snapshot copy also includes those blocks trapped by other Snapshot copies. On a Snapshot copy deletion, the "size" amount of blocks is the maximum number of blocks available. On a Snapshot copy restore, the "afs-used size" value will match the Snapshot copy "size" value.

|snaplock_expiry_time
|string

a|SnapLock expiry time for the Snapshot copy, if the Snapshot copy is taken on a SnapLock volume. A Snapshot copy is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time. This option can be set during Snapshot copy POST and Snapshot copy PATCH on Snapshot copy locking enabled volumes.

|snapmirror_label
|string

a|Label for SnapMirror operations

|state
|string

a|State of the Snapshot copy. There are cases where some Snapshot copies are not complete. In the "partial" state, the Snapshot copy is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial Snapshot copies cannot be used for a Snapshot copy restore operation. A Snapshot copy is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a Snapshot copy is valid.

|svm

```

|link:#svm[svm]
a|

|uuid
|string
a|The UUID of the Snapshot copy in the volume that uniquely identifies the
Snapshot copy in that volume.

|version_uuid
|string
a|The 128 bit identifier that uniquely identifies a snapshot and its
logical data layout.

|volume
|link:#volume[volume]
a|

|===

[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

[#error_arguments]
[.api-collapsible-fifth-title]

```

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

```

|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

:leveloffset: -1

= Retrieve clients

:leveloffset: +1

[[ID49039b2fe0cbdda6a5bc03e9e7be395c]]
= Storage volumes volume.uuid top-metrics clients endpoint overview

```

== Overview

You can use this API to retrieve a list of clients with the most IO activity for a specified volume. Use the `top_metric` parameter to specify which type of IO activity to filter for. This API is used to provide insight into IO activity and supports ordering by IO activity types, namely `iops` and `throughput` metrics. This API supports only returning one IO activity type per request.

== Failure to return list of clients with most IO activity

The API can sometimes fail to return the list of clients with the most IO activity, due to the following reasons:

- – The volume does not have the activity tracking feature enabled.

- – The volume does not have read/write traffic.

- – The read traffic is served by the NFS/CIFS client filesystem cache.

- – On rare occasions, the incoming traffic pattern is not suitable to

obtain the list of clients with the most IO activity.

== Retrieve a list of the clients with the most IO activity

For a report on the clients with the most IO activity returned in descending order, specify the IO activity type you want to filter for by passing the `iops` or `throughput` IO activity type into the `top_metric` parameter. If the IO activity type is not specified, by default the API returns a list of clients with the greatest number of average read operations per second. The maximum number of clients returned by the API for an IO activity type is 25.

– GET `/api/storage/volumes/{volume.uuid}/top-metrics/clients`

== Examples

=== Retrieving a list of the clients with the greatest average number of write operations per second:

The API:

```
GET /api/storage/volumes/{volume.uuid}/top-metrics/clients
```

The Call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-metrics/clients?top_metric=iops.write"
```

The Response:

```
{
  "records": [
    {
      "volume": {
        "name": "vol1"
      },
      "iops": {
        "write": 1495,
        "error": {
          "lower_bound": 1495,
          "upper_bound": 1505
        }
      },
      "client_ip": "172.28.71.128",
      "svm": {
        "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
        "name": "vs1",

```

```

    "_links": {
      "self": {
        "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
      }
    }
  },
  {
    "volume": {
      "name": "vol1"
    },
    "iops": {
      "write": 1022,
      "error": {
        "lower_bound": 1022,
        "upper_bound": 1032
      }
    },
    "client_ip": "172.28.71.179",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    }
  },
  {
    "volume": {
      "name": "vol1"
    },
    "iops": {
      "write": 345,
      "error": {
        "lower_bound": 345,
        "upper_bound": 355
      }
    },
    "client_ip": "172.28.51.62",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {

```

```

        "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/top-metrics/clients?top_metric=iops.write"
  }
}
}
}
-----

```

== Example showing the behavior of the API when there is no read/write traffic:

```

-----

# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/clients?top_metric=throughput.write"

# The Response:
{
"records": [
],
"num_records": 0,
"notice": {
  "message": "The activity tracking report for volume \"FV\" in SVM
\"vs0\" returned zero records. Check whether the volume have read/write
traffic. Refer to the REST API documentation for more information on why
there might be no records.",
  "code": "124518418"
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/9af63729-8ac8-11ec-b1bc-005056a79da4/top-metrics/clients?top_metric=throughput.write"
  }
}
}
}
-----

```



```
[[ID8809c6b6183556aef3173102cf835b12]]
```

```
= Retrieve clients with the most I/O activity
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/storage/volumes/{volume.uuid}/top-metrics/clients`#
```

```
*Introduced In:* 9.10
```

```
Retrieves a list of clients with the most IO activity.
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|volume.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Volume UUID
```

```
|top_metric
```

```
|string
```

```
|query
```

```
|False
```

```
a|IO activity type
```

```
* Default value: 1
```

```
* enum: ["iops.read", "iops.write", "throughput.read", "throughput.write"]
```

```
|svm.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by svm.uuid
```

```
|svm.name  
|string  
|query  
|False  
a|Filter by svm.name
```

```
|client_ip  
|string  
|query  
|False  
a|Filter by client_ip
```

```
|volume.name  
|string  
|query  
|False  
a|Filter by volume.name
```

```
|throughput.write  
|integer  
|query  
|False  
a|Filter by throughput.write
```

```
|throughput.error.lower_bound  
|integer  
|query  
|False  
a|Filter by throughput.error.lower_bound
```

```
|throughput.error.upper_bound  
|integer  
|query  
|False  
a|Filter by throughput.error.upper_bound
```

```
|throughput.read  
|integer  
|query
```

```
|False
a|Filter by throughput.read

|iops.read
|integer
|query
|False
a|Filter by iops.read

|iops.error.lower_bound
|integer
|query
|False
a|Filter by iops.error.lower_bound

|iops.error.upper_bound
|integer
|query
|False
a|Filter by iops.error.upper_bound

|iops.write
|integer
|query
|False
a|Filter by iops.write

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
```

```

|boolean
|query
|False
a|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
a|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
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===
== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```

|incomplete_response_reason
|link:#incomplete_response_reason[incomplete_response_reason]
a|Indicates that the metric report provides incomplete data.

|notice
|link:#notice[notice]
a|Optional field that indicates why no records are returned by the volume
activity tracking REST API.

|num_records
|integer
a|Number of records.

|records
|array[link:#top_metrics_client[top_metrics_client]]
a|

|===

```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "incomplete_response_reason": {
    "code": "111411207",
    "message": "Partial data has been returned for this metric report.
Reason: The activity tracking report for this volume is not available
because the system is busy collecting tracking data."
  },
  "notice": {
    "code": "111411207",
    "message": "No read/write traffic on volume"
  },
}

```

```
"num_records": 1,
"records": {
  "client_ip": "192.168.185.170",
  "iops": {
    "error": {
      "lower_bound": 34,
      "upper_bound": 54
    },
    "read": 5,
    "write": 10
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "throughput": {
    "error": {
      "lower_bound": 34,
      "upper_bound": 54
    },
    "read": 12,
    "write": 2
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
}
====

== Error
```

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 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.

| 124518415

| Failed to get the volume activity tracking report on volume volume.name in SVM svm.name.

| 124518416

| Internal error. Volume activity tracking report timed out for volume volume.name in SVM svm.name.

| 124518417

| Volume wildcard queries are not supported for activity tracking reports.

| 124518418

| The activity tracking report for volume volume.name in SVM svm.name returned zero records. Check whether the volume has read/write traffic.

Refer to the REST API documentation for more information.

```
| 124519413
| Volume activity tracking is not enabled on the volume.
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```
.Example error
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

```
====
```

```
== Definitions
```

```
[.api-def-first-level]
```

```
.See Definitions
```

```
[%collapsible%closed]
```

```
//Start collapsible Definitions block
```

```
====
```

```
[#href]
```

```
[.api-collapsible-fifth-title]
```



```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|next
```

```
|link:href[href]
```

```
a|
```

```
|self
```

```
|link:href[href]
```

```
a|
```

```
|===
```

```
[#incomplete_response_reason]
```

```
[.api-collapsible-fifth-title]
```

```
incomplete_response_reason
```

Indicates that the metric report provides incomplete data.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Warning code indicating why partial data was reported.
```

```
|message
```

```
|string
```

```
a|A message describing the reason for partial data.
```

```
|===
```

```
[#notice]
```

```
[.api-collapsible-fifth-title]
```

```
notice
```

Optional field that indicates why no records are returned by the volume activity tracking REST API.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Warning code indicating why no records are returned.
```

```
|message
```

```
|string
```

```
a|Details why no records are returned.
```

```
|===
```

```
[#top_metric_value_error_bounds]
```

```
[.api-collapsible-fifth-title]
```

```
top_metric_value_error_bounds
```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|lower_bound
|integer
a|Lower bound of the nominal value of a metric.

|upper_bound
|integer
a|Upper bound of the nominal value of a metric.

|===

[#iops]
[.api-collapsible-fifth-title]
iops

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#top_metric_value_error_bounds[top_metric_value_error_bounds]
a|

|read
|integer
a|Average number of read operations per second.

|write
|integer
a|Average number of write operations per second.

|===

[#_links]
[.api-collapsible-fifth-title]

```

`_links`

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the SVM.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the SVM.
```

```
|===
```

```
[#throughput]
```

```
[.api-collapsible-fifth-title]
```

```
throughput
```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|error
|link:#top_metric_value_error_bounds[top_metric_value_error_bounds]
a|

|read
|integer
a|Average number of read bytes received per second.

|write
|integer
a|Average number of write bytes received per second.

|===

[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume.

|uuid
|string
a|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
```

```
|===
```

```
[#top_metrics_client]
```

```
[.api-collapsible-fifth-title]
```

```
top_metrics_client
```

Information about a client's IO activity.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|client_ip
```

```
|string
```

a|IP address of the client. Both IPv4 and IPv6 IP addresses are supported.

```
|iops
```

```
|link:#iops[iops]
```

```
a|
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|throughput
```

```
|link:#throughput[throughput]
```

```
a|
```

```
|volume
```

```
|link:#volume[volume]
```

```
a|
```

```
|===
```

```
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
```

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

```

|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

:leveloffset: -1

= Retrieve directories

:leveloffset: +1

[[ID65112c0cdf04c6b112d0bcce372d30b6]]
= Storage volumes volume.uuid top-metrics directories endpoint overview

== Overview

You can use this API to retrieve a list of directories with the greatest
value performance metric or capacity metric for a specified volume. Use
the `top_metric` parameter to specify which type of metric to filter for.
This API is used to provide insight into IO metrics, namely `iops`,
`throughput`, and `non_recursive_bytes_used`. This API only supports
returning one metric per request.

== Retrieve a list of the directories with the most IO activity

For a report on the directories with the most IO activity returned in
descending order, specify the performance metric type you want to filter
for by passing the `iops` or `throughput` property into the top_metric
parameter. If the metric type is not specified, by default the API returns
a list of the directories with the greatest number of the average read
operations per second. The maximum number of directories returned by the
API for a metric type is 25.

== Retrieve a list of the largest directories

For a report on the largest directories returned in descending order,

```


specify the capacity metric by passing the `non_recursive_bytes_used` property into the `top_metric` parameter. If the metric type is not specified, by default the API returns a list of directories with the greatest number of average read operations per second. The maximum number of directories returned by the API for a metric type is 25.

== Failure to return list of directories with most IO activity

The API can sometimes fail to return the list of directories with the most IO activity, due to the following reasons:

– The volume does not have the activity tracking feature enabled.

– The volume does not have read/write traffic.

– The read traffic is served by the NFS/CIFS client filesystem cache.

– On rare occasions, the incoming traffic pattern is not suitable to obtain the list of directories with the most IO activity.

== Failure to return list of largest directories

The API can sometimes fail to return the list of largest directories, due to the following reasons:

– The volume does not have file system analytics enabled.

– The volume's file system analytics database version doesn't support this report.

== Failure to return the pathnames for the list of directories with most IO activity

The API can sometimes fail to obtain the filesystem pathnames for the list of directory entries, due to internal transient errors.

In such cases, instead of the pathname, the API will return

"{+++<volume_instance_uuid>+++:+<fileid>+++}" for every directory entry.

You can get more information about the directory entry by invoking the GET on the below API using the above obtained fileid. "GET

/api/storage/volumes/{+++<volume_instance_uuid>+++}/files/{path}?inode_number=+++<fileid>+++"+</fileid>++++++</volume_instance_uuid>++++++</fileid>++++++</volume_instance_uuid>+++"

== Examples

=== Retrieving a list of the directories with the greatest average number

of read operations per second

The API:

```
GET /api/storage/volumes/{volume.uuid}/top-metrics/directories
```

The Call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-  
metrics/directories?top_metric=iops.read"
```

The Response:

```
{  
  "records": [  
    {  
      "volume": {  
        "name": "voll1",  
      },  
      "iops": {  
        "read": 1495,  
        "error": {  
          "lower_bound": 1495,  
          "upper_bound": 1505  
        }  
      },  
      "path": "/dir1/dir2",  
      "svm": {  
        "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",  
        "name": "vs1",  
        "_links": {  
          "self": {  
            "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"  
          }  
        }  
      },  
      "_links": {  
        "directory": {  
          "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-  
2df8e52ef864/files/dir1%2Fdir2"  
        },  
        "metadata": {  
          "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-  
2df8e52ef864/files/dir1%2Fdir2?return_metadata=true"  
        }  
      }  
    },  
  ],  
}
```

```

{
  "volume": {
    "name": "vol1",
  },
  "iops": {
    "read": 1022,
    "error": {
      "lower_bound": 1022,
      "upper_bound": 1032
    }
  },
  "path": "/dir3/dir4",
  "svm": {
    "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
      }
    }
  },
  "_links": {
    "directory": {
      "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/files/dir3%2Fdir4"
    },
    "metadata": {
      "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/files/dir3%2Fdir4?return_metadata=true"
    }
  }
},
{
  "volume": {
    "uuid": "73b293df-e9d7-46cc-a9ce-2df8e52ef864",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864"
      }
    }
  },
  "iops": {
    "read": 345,
    "error": {

```

```

        "lower_bound": 345,
        "upper_bound": 355
    }
},
"path": "/dir12",
"svm": {
    "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
    "name": "vs1",
    "_links": {
        "self": {
            "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
    }
},
"_links": {
    "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/files/dir12"
    },
    "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/files/dir12?return_metadata=true"
    }
}
}
],
"num_records": 3,
"_links": {
    "self": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/top-metrics/directories?top_metric=iops.read"
    }
}
}
}
-----

```

=== Retrieving a list of the directories with the most traffic with failure in obtaining the pathnames for the directories:

The Call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-metrics/directories?top_metric=throughput.write"
```

The Response:

```

{
  "records": [
    {
      "volume": {
        "name": "fv"
      },
      "throughput": {
        "write": 24,
        "error": {
          "lower_bound": 24,
          "upper_bound": 29
        }
      },
      "path": "{4ec6d1ea-d5da-11eb-a25f-005056ac0f77:1232}",
      "svm": {
        "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
        "name": "vs0",
        "_links": {
          "self": {
            "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
          }
        }
      }
    },
    {
      "volume": {
        "name": "fv"
      },
      "throughput": {
        "write": 12,
        "error": {
          "lower_bound": 12,
          "upper_bound": 22
        }
      },
      "path": "{4ec6d1ea-d5da-11eb-a25f-005056ac0f77:6754}",
      "svm": {
        "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
        "name": "vs0",
        "_links": {
          "self": {
            "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
          }
        }
      }
    }
  ],
}

```

```

{
  "volume": {
    "name": "fv"
  },
  "throughput": {
    "write": 8,
    "error": {
      "lower_bound": 8,
      "upper_bound": 10
    }
  },
  "path": "{4ec6d1ea-d5da-11eb-a25f-005056ac0f77:8654}",
  "svm": {
    "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
    "name": "vs0",
    "_links": {
      "self": {
        "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
      }
    }
  }
},
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/4ec6d1ea-d5da-11eb-a25f-005056ac0f77/top-metrics/directories?top_metric=throughput.write"
  }
}
}

```

==== Retrieving a list of the largest directories

The following example shows how to retrieve a list of the largest directories based on the non-recursive bytes used by the contents of a directory.

The API:

```
GET /api/storage/volumes/{volume.uuid}/top-metrics/directories
```

The Call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
```

```
metrics/directories?top_metric=non_recursive_bytes_used"
```

```
# The Response:
```

```
{
"records": [
  {
    "volume": {
      "name": "vol1"
    },
    "non_recursive_bytes_used": 345,
    "path": "/dir1/dir2",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/files/dir1%2Fdir2"
      },
      "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/files/dir1%2Fdir2?return_metadata=true"
      }
    }
  },
  {
    "volume": {
      "name": "vol1"
    },
    "non_recursive_bytes_used": 345,
    "path": "/dir3/dir4",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    }
  },
}
```

```

    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/files/dir3%2Fdir4"
      },
      "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/files/dir3%2Fdir4?return_metadata=true"
      }
    },
    {
      "volume": {
        "name": "vol1"
      },
      "non_recursive_bytes_used": 345,
      "path": "/dir12",
      "svm": {
        "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
          }
        }
      },
      "_links": {
        "directory": {
          "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/files/dir12"
        },
        "metadata": {
          "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/files/dir12?return_metadata=true"
        }
      }
    }
  ],
  "num_records": 3,
  "_links": {
    "self": {
      "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/top-metrics/directories?top_metric=non_recursive_bytes_used"
    }
  }
}

```


==== Retrieving a list of the largest directories where incomplete data is reported

The following example retrieves a list of the largest directories when partial data is returned.

```
-----

# The API:
GET /api/storage/volumes/{volume.uuid}/top-metrics/directories

# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/directories?top_metric=non_recursive_bytes_used"

# The Response:
{
"records": [
  {
    "volume": {
      "name": "vol1"
    },
    "non_recursive_bytes_used": 1022,
    "path": "/dir1/dir2",
    "svm": {
      "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
        }
      }
    },
    "_links": {
      "directory": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir1%2Fdir2"
      },
      "metadata": {
        "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/files/dir1%2Fdir2?return_metadata=true"
      }
    }
  },
],
}
```

```

{
  "volume": {
    "name": "vol1"
  },
  "non_recursive_bytes_used": 261,
  "path": "/dir12",
  "svm": {
    "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
      }
    }
  },
  "_links": {
    "directory": {
      "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/files/dir12"
    },
    "metadata": {
      "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/files/dir12?return_metadata=true"
    }
  }
},
"num_records": 2,
"incomplete_response_reason": {
  "message": "Partial data has been returned for this metric report. Reason: Data collection for the large directory report is in progress.",
  "code": "111411234",
  "arguments": [
    {
      "message": "Data collection for the large directory report is in progress."
    }
  ]
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-2df8e52ef864/top-metrics/directories?top_metric=non_recursive_bytes_used"
  }
}
}

```

====
Retrieving a list of the largest directories when the file system analytics database version doesn't support this report

The following example shows the behavior of the API when the file system analytics database version doesn't support the large directory report.

====

```
# The API:
GET /api/storage/volumes/{volume.uuid}/top-metrics/directories

# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
metrics/directories?top_metric=non_recursive_bytes_used"

# The Response:
{
  "records": [
  ],
  "num_records": 0,
  "_links": {
    "self": {
      "href": "/api/storage/volumes/73b293df-e9d7-46cc-a9ce-
2df8e52ef864/top-metrics/directories?top_metric=non_recursive_bytes_used"
    }
  },
  "error": {
    "message": "The large directory report for volume \"FV\" in Vserver
\"vs0\" is not available because the file system analytics database
version doesn't support this report. Use the \"volume analytics off\"
command to disable analytics on the volume, then use the \"volume
analytics on\" command to re-enable analytics.",
    "code": "124519410"
  }
}
```

====
Example showing the behavior of the API when there is no read/write traffic:

====

```
# The Call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-
```

```
metrics/directories?top_metric=throughput.write"
```

```
# The Response:
```

```
{
"records": [
],
"num_records": 0,
"notice": {
  "message": "The activity tracking report for volume \"FV\" in SVM
\"vs0\" returned zero records. Check whether the volume have read/write
traffic. Refer to the REST API documentation for more information on why
there might be no records.",
  "code": "124518418"
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/9af63729-8ac8-11ec-b1bc-
005056a79da4/top-metrics/directories?top_metric=throughput.write"
  }
}
}
}
-----
```

```
[[ID11568dc672a1d174b310c85a64745a57]]
```

```
= Retrieve directories with the greatest value performance metric or
capacity metric
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/volumes/{volume.uuid}/top-metrics/directories`#
```

```
*Introduced In:* 9.10
```

```
Retrieves a list of directories with the greatest value performance metric
or capacity metric.
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```

|In
|Required
|Description

|volume.uuid
|string
|path
|True
a|Volume UUID

|top_metric
|string
|query
|False
a|Type of performance metric or capacity metric.

* Default value: 1
* enum: ["iops.read", "iops.write", "throughput.read", "throughput.write",
"non_recursive_bytes_used"]

|max_records_per_volume
|integer
|query
|False
a|Max records per volume.

|non_recursive_bytes_used
|integer
|query
|False
a|Filter by non_recursive_bytes_used

* Introduced in: 9.12

|throughput.read
|integer
|query
|False
a|Filter by throughput.read

|throughput.error.lower_bound
|integer

```

```
|query  
|False  
a|Filter by throughput.error.lower_bound
```

```
|throughput.error.upper_bound  
|integer  
|query  
|False  
a|Filter by throughput.error.upper_bound
```

```
|throughput.write  
|integer  
|query  
|False  
a|Filter by throughput.write
```

```
|svm.uuid  
|string  
|query  
|False  
a|Filter by svm.uuid
```

```
|svm.name  
|string  
|query  
|False  
a|Filter by svm.name
```

```
|iops.write  
|integer  
|query  
|False  
a|Filter by iops.write
```

```
|iops.error.lower_bound  
|integer  
|query  
|False  
a|Filter by iops.error.lower_bound
```

```
|iops.error.upper_bound
|integer
|query
|False
a|Filter by iops.error.upper_bound
```

```
|iops.read
|integer
|query
|False
a|Filter by iops.read
```

```
|path
|string
|query
|False
a|Filter by path
```

```
|volume.name
|string
|query
|False
a|Filter by volume.name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|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

a|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

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links

|link:#_links[_links]

a|

|incomplete_response_reason

|link:#incomplete_response_reason[incomplete_response_reason]

a|Indicates that the metric report provides incomplete data.


```
|notice
|link:#notice[notice]
a|Optional field that indicates why no records are returned by the volume
activity tracking REST API.
```

```
|num_records
|integer
a|Number of records.
```

```
|records
|array[link:#top_metrics_directory[top_metrics_directory]]
a|
```

```
|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "incomplete_response_reason": {
    "code": "111411207",
    "message": "Partial data has been returned for this metric report.
Reason: Data collection for the large directory report is in progress."
  },
  "notice": {
    "code": "111411207",
    "message": "No read/write traffic on volume."
  },
  "num_records": 1,
  "records": {
    "_links": {
      "metadata": {
        "href": "/api/resourcelink"
      }
    }
  }
}
```

```

    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "iops": {
    "error": {
      "lower_bound": 34,
      "upper_bound": 54
    },
    "read": 10,
    "write": 5
  },
  "non_recursive_bytes_used": 300,
  "path": "/dir_abc/dir_123/dir_20",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "throughput": {
    "error": {
      "lower_bound": 34,
      "upper_bound": 54
    },
    "read": 3,
    "write": 20
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 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.

| 124518415

| Failed to get the volume activity tracking report on volume volume.name in SVM svm.name.

| 124518416

| Internal error. Volume activity tracking report timed out for volume volume.name in SVM svm.name.

| 124518417

| Volume wildcard queries are not supported for activity tracking reports.

```
| 124518418
| The activity tracking report for volume volume.name in SVM svm.name
returned zero records. Check whether the volume has read/write traffic.
Refer to the REST API documentation for more information.
```

```
| 124519410
| The large directory report for volume volume.name in SVM svm.name is not
available because the file system analytics database version doesn't
support this report.
```

```
| 124519413
| Volume activity tracking is not enabled on the volume.
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```
.Example error
```

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

```
=====
```

```
== Definitions
```

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#incomplete_response_reason]
[.api-collapsible-fifth-title]
incomplete_response_reason

```

Indicates that the metric report provides incomplete data.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Warning code indicating why partial data was reported.
```

```
|message
```

```
|string
```

```
a|A message describing the reason for partial data.
```

```
|===
```

```
[#notice]
```

```
[.api-collapsible-fifth-title]
```

```
notice
```

Optional field that indicates why no records are returned by the volume activity tracking REST API.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Warning code indicating why no records are returned.
```

```
|message
```

```
|string
```

```
a|Details why no records are returned.
```

```

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|metadata
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#top_metric_value_error_bounds]
[.api-collapsible-fifth-title]
top_metric_value_error_bounds

[cols=3*,options=header]
|===
|Name
|Type
|Description

|lower_bound
|integer
a|Lower bound of the nominal value of a metric.

|upper_bound
|integer
a|Upper bound of the nominal value of a metric.

|===

```

```

[#iops]
[.api-collapsible-fifth-title]
iops

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#top_metric_value_error_bounds[top_metric_value_error_bounds]
a|

|read
|integer
a|Average number of read operations per second.

|write
|integer
a|Average number of write operations per second.

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

```



```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#top_metric_value_error_bounds[top_metric_value_error_bounds]
a|

|read
|integer
a|Average number of read bytes received per second.

|write
|integer
a|Average number of write bytes received per second.

```

```
|===
```

```
[#volume]  
[.api-collapsible-fifth-title]  
volume
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|name  
|string  
a|The name of the volume.
```

```
|uuid  
|string  
a|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
```

```
|===
```

```
[#top_metrics_directory]  
[.api-collapsible-fifth-title]  
top_metrics_directory
```

```
Information about a directory's IO metrics.
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```

|Type
|Description

|_links
|link:#_links[_links]
a|

|iops
|link:#iops[iops]
a|

|non_recursive_bytes_used
|integer
a|Non-recursive bytes used by the contents of a directory.

|path
|string
a|Path of the directory.

|svm
|link:#svm[svm]
a|

|throughput
|link:#throughput[throughput]
a|

|volume
|link:#volume[volume]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code

```

```

|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
:leveloffset: -1
```

```
= Retrieve files
```

```
:leveloffset: +1
```

```
[[ID9f33b1b0a937a6d732075ce6ceed7981]]
```

```
= Storage volumes volume.uuid top-metrics files endpoint overview
```

```
== Overview
```

You can use this API to retrieve a list of files with the most IO activity for a specified volume. Use the `top_metric` parameter to specify which type of IO activity to filter for. This API is used to provide insight into IO activity and supports ordering by IO activity types, namely `iops` or `throughput` metrics. This API also supports only returning one IO activity type per request.

```
== Failure to return list of files with most IO activity
```

The API can sometimes fail to return the list of files with the most IO activity, due to the following reasons:

– The volume does not have the activity tracking feature enabled.

– The volume does not have read/write traffic.

– The read traffic is served by the NFS/CIFS client filesystem cache.

– On rare occasions, the incoming traffic pattern is not suitable to obtain the list of files with the most IO activity.

```
== Failure to return the pathnames for the list of files with most IO activity
```

The API can sometimes fail to obtain the filesystem pathnames for the list of files, due to internal transient errors.

In such cases, instead of the pathname, the API will return

"{+++<volume_instance_uuid>+++::+++<fileid>+++}" for every file entry. You

can get more information about the file entry by invoking the GET on the below API using the above obtained fileid. "GET /api/storage/volumes/{+++<volume_instance_uuid>+++}/files/{path}?inode_number=+++<fileid>+++"+++</fileid>++++++</volume_instance_uuid>++++++</fileid>++++++</volume_instance_uuid>+++

== Retrieve a list of the files with the most IO activity

For a report on the files with the most IO activity returned in descending order, specify the IO activity type you want to filter for by passing the `iops` or `throughput` property into the top_metric parameter. If the IO activity type is not specified, by default the API returns a list of the files with the greatest number of the average read operations per second. The maximum number of files returned by the API for an IO activity type is 25.

– GET /api/storage/volumes/{volume.uuid}/top-metrics/files

== Examples

=== Retrieving a list of the files with the greatest average number of write bytes received per second:

The API:

GET /api/storage/volumes/{volume.uuid}/top-metrics/files

The Call:

curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-metrics/files?top_metric=throughput.write"

The Response:

```
{
  "records": [
    {
      "volume": {
        "name": "fv"
      },
      "throughput": {
        "write": 24,
        "error": {
          "lower_bound": 24,
          "upper_bound": 29
        }
      },
      "path": "/d5/f5",
```

```

"svm": {
  "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
  "name": "vs0",
  "_links": {
    "self": {
      "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
    }
  }
},
"_links": {
  "metadata": {
    "href": "/api/storage/volumes/4ec6d1ea-d5da-11eb-a25f-005056ac0f77/files/d5%2Ff5?return_metadata=true"
  }
}
},
{
  "volume": {
    "name": "fv"
  },
  "throughput": {
    "write": 12,
    "error": {
      "lower_bound": 12,
      "upper_bound": 22
    }
  },
  "path": "/d6/f6",
  "svm": {
    "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
    "name": "vs0",
    "_links": {
      "self": {
        "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
      }
    }
  }
},
"_links": {
  "metadata": {
    "href": "/api/storage/volumes/4ec6d1ea-d5da-11eb-a25f-005056ac0f77/files/d6%2Ff6?return_metadata=true"
  }
}
},
{
  "volume": {

```

```

    "name": "fv"
  },
  "throughput": {
    "write": 8,
    "error": {
      "lower_bound": 8,
      "upper_bound": 10
    }
  },
  "path": "/d3/f3",
  "svm": {
    "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
    "name": "vs0",
    "_links": {
      "self": {
        "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
      }
    }
  },
  "_links": {
    "metadata": {
      "href": "/api/storage/volumes/4ec6d1ea-d5da-11eb-a25f-005056ac0f77/files/d3%2Ff3?return_metadata=true"
    }
  }
},
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/4ec6d1ea-d5da-11eb-a25f-005056ac0f77/top-metrics/files?top_metric=throughput.write"
  }
}
}
}
-----

```

=== Retrieving a list of the files with the most traffic with failure in obtaining the pathnames for the files:

The Call:

```

curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-metrics/files?top_metric=throughput.write"

```



```

# The Response:
{
  "records": [
    {
      "volume": {
        "name": "fv"
      },
      "throughput": {
        "write": 24,
        "error": {
          "lower_bound": 24,
          "upper_bound": 29
        }
      },
      "path": "{4ec6d1ea-d5da-11eb-a25f-005056ac0f77:1232}",
      "svm": {
        "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
        "name": "vs0",
        "_links": {
          "self": {
            "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
          }
        }
      }
    },
    {
      "volume": {
        "name": "fv"
      },
      "throughput": {
        "write": 12,
        "error": {
          "lower_bound": 12,
          "upper_bound": 22
        }
      },
      "path": "{4ec6d1ea-d5da-11eb-a25f-005056ac0f77:6754}",
      "svm": {
        "uuid": "0ba74c3e-d5ca-11eb-8fbb-005056ac0f77",
        "name": "vs0",
        "_links": {
          "self": {
            "href": "/api/svm/svms/0ba74c3e-d5ca-11eb-8fbb-005056ac0f77"
          }
        }
      }
    }
  ]
}

```



```

],
"num_records": 0,
"notice": {
  "message": "The activity tracking report for volume \"FV\" in SVM
\"vs0\" returned zero records. Check whether the volume have read/write
traffic. Refer to the REST API documentation for more information on why
there might be no records.",
  "code": "124518418"
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/9af63729-8ac8-11ec-b1bc-
005056a79da4/top-metrics/files?top_metric=throughput.write"
  }
}
}
}
-----

```

```

[[IDf526d22a8b322630def994ad33f464d3]]
= Retrieve files with the most I/O activity

```

```

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/volumes/{volume.uuid}/top-metrics/files`#

```

Introduced In: 9.10

Retrieves a list of files with the most IO activity.

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```

|Name
|Type
|In
|Required
|Description

```

```

|volume.uuid
|string
|path

```

```
|True
a|Volume UUID

|top_metric
|string
|query
|False
a|IO activity type

* Default value: 1
* enum: ["iops.read", "iops.write", "throughput.read", "throughput.write"]

|max_records_per_volume
|integer
|query
|False
a|Max records per volume.

|iops.read
|integer
|query
|False
a|Filter by iops.read

|iops.error.lower_bound
|integer
|query
|False
a|Filter by iops.error.lower_bound

|iops.error.upper_bound
|integer
|query
|False
a|Filter by iops.error.upper_bound

|iops.write
|integer
|query
|False
a|Filter by iops.write
```

```
|volume.name  
|string  
|query  
|False  
a|Filter by volume.name
```

```
|path  
|string  
|query  
|False  
a|Filter by path
```

```
|throughput.read  
|integer  
|query  
|False  
a|Filter by throughput.read
```

```
|throughput.error.lower_bound  
|integer  
|query  
|False  
a|Filter by throughput.error.lower_bound
```

```
|throughput.error.upper_bound  
|integer  
|query  
|False  
a|Filter by throughput.error.upper_bound
```

```
|throughput.write  
|integer  
|query  
|False  
a|Filter by throughput.write
```

```
|svm.uuid  
|string  
|query
```

```
|False
a|Filter by svm.uuid

|svm.name
|string
|query
|False
a|Filter by svm.name

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|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
a|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
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|incomplete_response_reason
|link:#incomplete_response_reason[incomplete_response_reason]
a|Indicates that the metric report provides incomplete data.

|notice
|link:#notice[notice]
a|Optional field that indicates why no records are returned by the volume
activity tracking REST API.

|num_records
|integer
a|Number of records.

|records
|array[link:#top_metrics_file[top_metrics_file]]
a|

|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "incomplete_response_reason": {
    "code": "111411207",
    "message": "Partial data has been returned for this metric report. Reason: The activity tracking report for this volume is not available because the system is busy collecting tracking data."
  },
  "notice": {
    "code": "111411207",
    "message": "No read/write traffic on volume."
  },
  "num_records": 1,
  "records": {
    "_links": {
      "metadata": {
        "href": "/api/resourcelink"
      },
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "iops": {
      "error": {
        "lower_bound": 34,
        "upper_bound": 54
      },
      "read": 5,
      "write": 4
    },
    "path": "/dir_abc/dir_123/file_1",
    "svm": {
      "_links": {
```



```

    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"throughput": {
  "error": {
    "lower_bound": 34,
    "upper_bound": 54
  },
  "read": 2,
  "write": 20
},
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 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.

| 124518415
| Failed to get the volume activity tracking report on volume volume.name
in SVM svm.name.

| 124518416
| Internal error. Volume activity tracking report timed out for volume
volume.name in SVM svm.name.

| 124518417
| Volume wildcard queries are not supported for activity tracking reports.

| 124518418
| The activity tracking report for volume volume.name in SVM svm.name
returned zero records. Check whether the volume has read/write traffic.
Refer to the REST API documentation for more information.

| 124519413
| Volume activity tracking is not enabled on the volume.
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|next
|link:#href[href]
a|
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#incomplete_response_reason]
[.api-collapsible-fifth-title]
incomplete_response_reason
```

Indicates that the metric report provides incomplete data.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Warning code indicating why partial data was reported.
```

```
|message
|string
a|A message describing the reason for partial data.
```

```
|===
```

```
[#notice]
[.api-collapsible-fifth-title]
notice
```

Optional field that indicates why no records are returned by the volume activity tracking REST API.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Warning code indicating why no records are returned.
```

```
|message
|string
a|Details why no records are returned.
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|metadata
|link:#href[href]
a|
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```

[#top_metric_value_error_bounds]
[.api-collapsible-fifth-title]
top_metric_value_error_bounds

[cols=3*,options=header]
|===
|Name
|Type
|Description

|lower_bound
|integer
a|Lower bound of the nominal value of a metric.

|upper_bound
|integer
a|Upper bound of the nominal value of a metric.

|===

[#iops]
[.api-collapsible-fifth-title]
iops

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#top_metric_value_error_bounds[top_metric_value_error_bounds]
a|

|read
|integer
a|Average number of read operations per second.

|write
|integer
a|Average number of write operations per second.

```

```

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

```

```

[#throughput]
[.api-collapsible-fifth-title]
throughput

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#top_metric_value_error_bounds[top_metric_value_error_bounds]
a|

|read
|integer
a|Average number of read bytes received per second.

|write
|integer
a|Average number of write bytes received per second.

|===

[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume.

```



```
|uuid
|string
a|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
```

```
|===
```

```
[#top_metrics_file]
[.api-collapsible-fifth-title]
top_metrics_file
```

Information about a file's IO activity.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|iops
|link:#iops[iops]
a|
```

```
|path
|string
a|Path of the file.
```

```
|svm
|link:#svm[svm]
a|
```

```
|throughput
|link:#throughput[throughput]
```

```

a|

|volume
|link:#volume[volume]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

```

```

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

:leveloffset: -1

= Retrieve users

:leveloffset: +1

[[ID507c6b7ec656fe9f0f039df791a76183]]
= Storage volumes volume.uuid top-metrics users endpoint overview

== Overview

You can use this API to retrieve a list of users with the most IO activity
for a specified volume. Use the `top_metric` parameter to specify which
type of IO activity to filter for. This API is used to provide insight
into IO activity and supports ordering by IO activity types, namely `iops`
or `throughput` metrics. This API also supports only returning one IO
activity type per request.

== Failure to return list of users with most IO activity

```

The API can sometimes fail to return the list of users with the most IO activity, due to the following reasons:

– The volume does not have the activity tracking feature enabled.

– The volume does not have read/write traffic.

– The read traffic is served by the NFS/CIFS client filesystem cache.

– On rare occasions, the incoming traffic pattern is not suitable to obtain the list of users with the most IO activity.

== Failure to return the usernames

The API can sometimes fail to obtain the usernames for the list of userid entries, due to internal transient errors.

In such cases, instead of the username, the API will return "{+++<user-id>+++}" for every user entry.+++</user-id>+++

== Retrieve a list of the users with the most IO activity

For a report on the users with the most IO activity returned in descending order, specify the IO activity type you want to filter for by passing the `iops` or `throughput` property into the `top_metric` parameter. If the IO activity type is not specified, by default the API returns a list of the users with the greatest number of the average read operations per second. The maximum number of users returned by the API for an IO activity type is 25.

– GET `/api/storage/volumes/{volume.uuid}/top-metrics/users`

== Examples

=== Retrieving a list of the users with the greatest average number of read bytes received per second:

The API:

```
GET /api/storage/volumes/{volume.uuid}/top-metrics/users
```

The Call:

```
curl -X GET "https://<mgmt-ip>/api/storage/volumes/{volume.uuid}/top-metrics/users?top_metric=throughput.read"
```

The Response:

```

{
  "records": [
    {
      "volume": {
        "name": "vol1"
      },
      "throughput": {
        "read": 1495,
        "error": {
          "lower_bound": 1495,
          "upper_bound": 1502
        }
      },
      "user_id": "S-1-5-21-256008430-3394229847-3930036330-1001",
      "user_name": "John",
      "svm": {
        "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
          }
        }
      }
    },
    {
      "volume": {
        "name": "vol1"
      },
      "throughput": {
        "read": 1022,
        "error": {
          "lower_bound": 1022,
          "upper_bound": 1025
        }
      },
      "user_id": "1988",
      "user_name": "Ryan",
      "svm": {
        "uuid": "572361f3-e769-439d-9c04-2ba48a08ff43",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/572361f3-e769-439d-9c04-2ba48a08ff43"
          }
        }
      }
    }
  ]
}

```



```

{
  "records": [
  ],
  "num_records": 0,
  "notice": {
    "message": "The activity tracking report for volume \"FV\" in SVM
\"vs0\" returned zero records. Check whether the volume have read/write
traffic. Refer to the REST API documentation for more information on why
there might be no records.",
    "code": "124518418"
  },
  "_links": {
    "self": {
      "href": "/api/storage/volumes/9af63729-8ac8-11ec-b1bc-
005056a79da4/top-metrics/users?top_metric=throughput.write"
    }
  }
}
}
}
----

```

```

[[ID4d16f046e06ea841f145635e9661c816]]
= Retrieve users with the most I/O activity

```

```

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/volumes/{volume.uuid}/top-metrics/users`#

```

Introduced In: 9.10

Retrieves a list of users with the most IO activity.

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|volume.uuid
```

```
|string
|path
|True
a|Volume UUID
```

```
|top_metric
|string
|query
|False
a|IO activity type
```

```
* Default value: 1
* enum: ["iops.read", "iops.write", "throughput.read", "throughput.write"]
```

```
|user_id
|string
|query
|False
a|Filter by user_id
```

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|throughput.write
|integer
|query
|False
a|Filter by throughput.write
```

```
|throughput.error.lower_bound
|integer
|query
```



```
|False
a|Filter by throughput.error.lower_bound

|throughput.error.upper_bound
|integer
|query
|False
a|Filter by throughput.error.upper_bound

|throughput.read
|integer
|query
|False
a|Filter by throughput.read

|volume.name
|string
|query
|False
a|Filter by volume.name

|user_name
|string
|query
|False
a|Filter by user_name

|iops.write
|integer
|query
|False
a|Filter by iops.write

|iops.read
|integer
|query
|False
a|Filter by iops.read

|iops.error.lower_bound
```

```
|integer
|query
|False
a|Filter by iops.error.lower_bound
```

```
|iops.error.upper_bound
|integer
|query
|False
a|Filter by iops.error.upper_bound
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|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
a|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

a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|incomplete_response_reason

|link:#incomplete_response_reason[incomplete_response_reason]

a|Indicates that the metric report provides incomplete data.

|notice

|link:#notice[notice]

a|Optional field that indicates why no records are returned by the volume
activity tracking REST API.

|num_records

|integer

a|Number of records.

|records

|array[link:#top_metrics_user[top_metrics_user]]

a|

```
|===
```

```
.Example response
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs+=macros]
```

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "incomplete_response_reason": {
    "code": "111411207",
    "message": "Partial data has been returned for this metric report. Reason: The activity tracking report for this volume is not available because the system is busy collecting tracking data."
  },
  "notice": {
    "code": "111411207",
    "message": "No read/write traffic on volume."
  },
  "num_records": 1,
  "records": {
    "iops": {
      "error": {
        "lower_bound": 34,
        "upper_bound": 54
      },
      "read": 4,
      "write": 8
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
}
```

```

"throughput": {
  "error": {
    "lower_bound": 34,
    "upper_bound": 54
  },
  "read": 10,
  "write": 7
},
"user_id": "S-1-5-21-256008430-3394229847-3930036330-1001",
"user_name": "James",
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|====
| Error Code | Description
| 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.

| 124518415

| Failed to get the volume activity tracking report on volume volume.name in SVM svm.name.

| 124518416

| Internal error. Volume activity tracking report timed out for volume volume.name in SVM svm.name.

| 124518417

| Volume wildcard queries are not supported for activity tracking reports.

| 124518418

| The activity tracking report for volume volume.name in SVM svm.name returned zero records. Check whether the volume has read/write traffic. Refer to the REST API documentation for more information.

| 124519413

| Volume activity tracking is not enabled on the volume.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|next
```

```
|link:href[href]
```

```
a|
```

```
|self
```

```
|link:href[href]
```

```
a|
```

```
|===
```

```
[#incomplete_response_reason]
```

```
[.api-collapsible-fifth-title]
```

```
incomplete_response_reason
```

Indicates that the metric report provides incomplete data.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Warning code indicating why partial data was reported.
```

```
|message
```

```
|string
```

```
a|A message describing the reason for partial data.
```

```
|===
```

```
[#notice]
```

```
[.api-collapsible-fifth-title]
```

```
notice
```


Optional field that indicates why no records are returned by the volume activity tracking REST API.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Warning code indicating why no records are returned.
```

```
|message
```

```
|string
```

```
a|Details why no records are returned.
```

```
|===
```

```
[#top_metric_value_error_bounds]
```

```
[.api-collapsible-fifth-title]
```

```
top_metric_value_error_bounds
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|lower_bound
```

```
|integer
```

```
a|Lower bound of the nominal value of a metric.
```

```
|upper_bound
```

```
|integer
```

```
a|Upper bound of the nominal value of a metric.
```

```
|===
```

```
[#iops]
```

```

[.api-collapsible-fifth-title]
iops

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#top_metric_value_error_bounds[top_metric_value_error_bounds]
a|

|read
|integer
a|Average number of read operations per second.

|write
|integer
a|Average number of write operations per second.

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#top_metric_value_error_bounds[top_metric_value_error_bounds]
a|

|read
|integer
a|Average number of read bytes received per second.

|write
|integer
a|Average number of write bytes received per second.

```

```

|===

[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume.

|uuid
|string
a|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

|===

[#top_metrics_user]
[.api-collapsible-fifth-title]
top_metrics_user

Information about a user's IO activity.

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|iops
|link:#iops[iops]
a|

|svm
|link:#svm[svm]
a|

|throughput
|link:#throughput[throughput]
a|

|user_id
|string
a|User ID of the user.

|user_name
|string
a|Name of the user.

|volume
|link:#volume[volume]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message

```

```
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1
```

:leveloffset: -1

:leveloffset: -1

<<<

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b) (3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement

is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at [link:http://www.netapp.com/TM](http://www.netapp.com/TM)^[http://www.netapp.com/TM^] are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.