



Manage storage aggregates

ONTAP 9.9.1 REST API reference

NetApp
August 29, 2024

Table of Contents

- Manage storage aggregates 1
 - Storage aggregates UUID endpoint overview 1
 - Delete an aggregate specified by the UUID 11
 - Retrieve an aggregate specified by the UUID 15
 - Update an aggregate specified by the UUID 35

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.
- 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.
- data_encryption.software_encryption_enabled - This property enables or disables NAE on the 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.

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
    },
    "cloud_storage": {
      "used": 0
    },
    "efficiency": {
      "savings": 1408029,
      "ratio": 6.908119720880661,
      "logical_used": 1646350
    },
    "efficiency_without_snapshots": {
      "savings": 0,
      "ratio": 1,
      "logical_used": 737280
    },
  },
}
```

```

"efficiency_without_snapshots_flexclones": {
  "savings": 5000,
  "ratio": 2,
  "logical_used": 10000
}
},
"state": "online",
"snaplock_type": "non_snaplock",
"create_time": "2018-12-04T15:40:38-05:00",
"data_encryption": {
  "software_encryption_enabled": false,
  "drive_protection_enabled": false
},
"block_storage": {
  "primary": {
    "disk_count": 6,
    "disk_class": "solid_state",
    "raid_type": "raid_dp",
    "raid_size": 24,
    "checksum_style": "block",
    "disk_type": "ssd"
  },
  "hybrid_cache": {
    "enabled": false
  },
  "mirror": {
    "enabled": false,
    "state": "unmirrored"
  },
  "plexes": [
    {
      "name": "plex0",
    }
  ]
},
"cloud_storage": {
  "attach_eligible": false
},
}

```

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\": 14}}}"

# The response:
{
  "warnings": [
    {
      "name": "node_2_SSD_1",
      "warning": {
        "message": "One or more disks will not be added. 14 disks specified, 13 disks will be added.",
        "code": 787170,
        "arguments": [
          "14",
          "13"
        ]
      }
    }
  ]
}

```

```

    }
  ],
  "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",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/4046dda8-f802-11e8-8f6d-005056bb2030"
          }
        }
      },
      "space": {
        "block_storage": {
          "size": 1116180480
        }
      },
      "block_storage": {
        "primary": {
          "disk_count": 23,
          "disk_class": "solid_state",
          "raid_type": "raid_dp",
          "disk_type": "ssd"
        },
        "hybrid_cache": {
          "enabled": false
        },
        "mirror": {
          "enabled": false
        }
      },
    }
  ],
  "job": {
    "_links": {
      "self": {
        "href": "/api/storage/aggregates/cae60cfe-deae-42bd-babb-ef437d118314"
      }
    }
  }
}

```



```
}
```

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",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/4046dda8-f802-11e8-8f6d-005056bb2030"
      }
    }
  },
  "space": {
    "block_storage": {
      "size": 1116180480
    }
  },
  "block_storage": {
    "primary": {
      "disk_count": 23,
      "disk_class": "solid_state",
      "raid_type": "raid_dp",
      "disk_type": "ssd"
    },
    "hybrid_cache": {
      "enabled": false
    }
  }
}
```

```

    },
    "mirror": {
      "enabled": false
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/aggregates/cae60cfe-deae-42bd-babb-ef437d118314"
    }
  }
}

```

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
    }
  },
}

```

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-be35-4d93-9590-5a402bffb4b" -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-be35-4d93-9590-5a402bfffbe4b?fields=data_encryption.software_encryption_enabled" -H
"accept: application/json"

# The response:
{
  "uuid": "f3aafdc6-be35-4d93-9590-5a402bfffbe4b",
  "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 |

| 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 |
|------------|--|
| 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 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.footprint`
- `statistics.*`

Related ONTAP commands

- `storage aggregate show`

Parameters

| Name | Type | In | Required | Description |
|------------------------------------|---------------|-------|----------|---|
| <code>uuid</code> | string | path | True | Aggregate UUID |
| <code>auto_provision_policy</code> | string | query | False | If set to expand, a query is run on the system for the recommended optimal expansion layout of the aggregate. • Introduced in: 9.8 |
| <code>fields</code> | array[string] | query | False | Specify the fields to return. |

Response

Status: 200, Ok

| Name | Type | Description |
|---------------------|------------------------|-------------|
| <code>_links</code> | _links | |

| Name | Type | Description |
|-------------------------|---|--|
| block_storage | block_storage | Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier. |
| cloud_storage | cloud_storage | Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier. |
| create_time | string | Timestamp of aggregate creation |
| data_encryption | data_encryption | |
| dr_home_node | dr_home_node | Node where the aggregate 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 | |
| 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. |
| snaplock_type | string | SnapLock type |
| space | space | |
| state | string | Operational state of the aggregate |
| statistics | statistics | The real time I/O statistics for the aggregate. |
| uuid | string | Aggregate UUID |

Example response

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

```

    },
    "name": "store1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "used": 0
}
]
},
"create_time": "2018-01-01T12:00:00-04:00",
"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": "2019-12-12T12:00:00-04:00"
},
"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": "2017-01-25T11:20:13Z"
},
"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"snaplock_type": "string",
"space": {
  "block_storage": {
    "available": "10156560384",
    "full_threshold_percent": 0,
    "inactive_user_data": "304448",
    "physical_used": "2461696",
    "size": "10156769280",
    "used": "2088960"
  },
  "cloud_storage": {
    "used": "402743264"
  },
  "efficiency": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots_flexclones": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "footprint": "608896"
},
"state": "string",
"statistics": {

```

```

    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "uuid": "string"
}

```

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. |

| Error Code | Description |
|------------|---|
| 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. |

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

Example error

```

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

```

Definitions

See Definitions

href

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

_links

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

hybrid_cache

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

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

mirror

| Name | Type | Description |
|---------|---------|-----------------------------------|
| enabled | boolean | Aggregate is SyncMirror protected |
| state | string | |

plex_reference

Plex

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

primary

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

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

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 | |

| Name | Type | Description |
|---------|-----------------------|---|
| plexes | array[plex_reference] | Plex reference for each plex in the aggregate. |
| primary | primary | Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details. |

cloud_store

Cloud store

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

cloud_storage_tier

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

cloud_storage

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

| Name | Type | Description |
|-----------------|---------------------------|---|
| attach_eligible | boolean | Specifies whether the aggregate is eligible for a cloud store to be attached. |
| stores | array[cloud_storage_tier] | Configuration information for each cloud storage portion of the aggregate. |

| 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. |

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. |

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 | |

block_storage

| Name | Type | Description |
|-----------|---------|---------------------------|
| available | integer | Space available in bytes. |

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

cloud_storage

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

efficiency

Storage efficiency

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

efficiency_without_snapshots

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

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

efficiency_without_snapshots_flexclones

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) |

space

| Name | Type | Description |
|---------------|-------------------------------|--------------------|
| block_storage | block_storage | |
| cloud_storage | cloud_storage | |
| efficiency | efficiency | Storage efficiency |

| Name | Type | Description |
|--|--|---|
| 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 FlexClones. |
| footprint | integer | A summation of volume footprints (including volume guarantees), in bytes. This includes all of the volume footprints in the block_storage tier and the cloud_storage tier. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** . |

iops_raw

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

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

latency_raw

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

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

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 | |
| 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. |

| Name | Type | Description |
|-------------------------|---|--|
| 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 | |
| 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. |
| snaplock_type | string | SnapLock type |
| space | space | |
| state | string | Operational state of the aggregate |
| statistics | statistics | The real time I/O statistics for the aggregate. |
| uuid | string | Aggregate UUID |

Example request

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

```

    },
    "name": "store1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "used": 0
}
]
},
"create_time": "2018-01-01T12:00:00-04:00",
"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": "2019-12-12T12:00:00-04:00"
},
"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": "2017-01-25T11:20:13Z"
},
"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"snaplock_type": "string",
"space": {
  "block_storage": {
    "available": "10156560384",
    "full_threshold_percent": 0,
    "inactive_user_data": "304448",
    "physical_used": "2461696",
    "size": "10156769280",
    "used": "2088960"
  },
  "cloud_storage": {
    "used": "402743264"
  },
  "efficiency": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots_flexclones": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "footprint": "608896"
},
"state": "string",
"statistics": {

```

```

    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "uuid": "string"
}

```

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

```

    }
  },
  "cloud_storage": {
    "stores": [
      {
        "cloud_store": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "store1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "used": 0
      }
    ]
  },
  "create_time": "2018-01-01T12:00:00-04:00",
  "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": "2019-12-12T12:00:00-04:00"
  },
  "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": "2017-01-25T11:20:13Z"
  },
  "name": "node1_aggr_1",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snaplock_type": "string",
  "space": {
    "block_storage": {
      "available": "10156560384",
      "full_threshold_percent": 0,
      "inactive_user_data": "304448",
      "physical_used": "2461696",
      "size": "10156769280",
      "used": "2088960"
    },
    "cloud_storage": {
      "used": "402743264"
    },
    "efficiency": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "efficiency_without_snapshots": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    }
  }
}

```

```

    },
    "efficiency_without_snapshots_flexclones": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "footprint": "608896"
  },
  "state": "string",
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "uuid": "string"
}
],
"warnings": [
  {
    "action": {
      "arguments": [
        "string"
      ],
      "message": "string"
    },
    "name": "string",
    "warning": {
      "arguments": [
        "string"
      ],
      "message": "string"
    }
  }
]

```



```
}  
]  
}
```

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. |
| 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. |

| Error Code | Description |
|-------------------|--|
| 786955 | Modifying raidtype to raid_tec requires a minimum of six disks in the RAID Group. |
| 786956 | Modifying raidtype to raid_dp requires a minimum of four disks in the RAID Group. |
| 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. |
| 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. |
| 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. |
| 13108106 | Cannot run aggregate relocation because volume expand is in progress. |

| Error Code | Description |
|-------------------|--|
| 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. |
| 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 | |

hybrid_cache

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

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

mirror

| Name | Type | Description |
|---------|---------|-----------------------------------|
| enabled | boolean | Aggregate is SyncMirror protected |
| state | string | |

plex_reference

Plex

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

primary

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

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

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 | |

| Name | Type | Description |
|---------|-----------------------|---|
| plexes | array[plex_reference] | Plex reference for each plex in the aggregate. |
| primary | primary | Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details. |

cloud_store

Cloud store

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

cloud_storage_tier

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

cloud_storage

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

| Name | Type | Description |
|-----------------|---------------------------|---|
| attach_eligible | boolean | Specifies whether the aggregate is eligible for a cloud store to be attached. |
| stores | array[cloud_storage_tier] | Configuration information for each cloud storage portion of the aggregate. |

| 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. |

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. |

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 | |

block_storage

| Name | Type | Description |
|-----------|---------|---------------------------|
| available | integer | Space available in bytes. |

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

cloud_storage

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

efficiency

Storage efficiency

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

efficiency_without_snapshots

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

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

efficiency_without_snapshots_flexclones

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) |

space

| Name | Type | Description |
|---------------|-------------------------------|--------------------|
| block_storage | block_storage | |
| cloud_storage | cloud_storage | |
| efficiency | efficiency | Storage efficiency |

| Name | Type | Description |
|--|--|---|
| 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 FlexClones. |
| footprint | integer | A summation of volume footprints (including volume guarantees), in bytes. This includes all of the volume footprints in the block_storage tier and the cloud_storage tier. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or ** . |

iops_raw

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

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

latency_raw

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

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

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 | |
| 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 | |
| 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. |
| snaplock_type | string | SnapLock type |
| space | space | |
| state | string | Operational state of the aggregate |
| statistics | statistics | The real time I/O statistics for the aggregate. |

| Name | Type | Description |
|------|--------|----------------|
| uuid | string | Aggregate UUID |

job_link

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

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. |

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

Copyright information

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

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

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

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

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

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

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

Trademark information

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