



Retrieve or create a collection of storage aggregates

ONTAP 9.9.1 REST API reference

NetApp
April 02, 2024

Table of Contents

- Retrieve or create a collection of storage aggregates 1
 - Storage aggregates endpoint overview 1
 - Retrieve a collection of aggregates for an entire cluster 6
 - Create a collection of aggregates for an entire cluster 39

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_spares". Using the "recommend" query returns the list of aggregates that are recommended for creation in the cluster. The "show_spares" 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.

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" and "space.block_storage.inactive_user_data" 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.
-

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": 24,
          "disk_class": "solid_state",
          "raid_type": "raid_dp",
          "disk_type": "ssd"
        },
        "hybrid_cache": {
          "enabled": false
        },
        "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": 10,
          "disk_class": "solid_state",

```

```

    "raid_type": "raid_dp",
    "disk_type": "ssd"
  },
  "hybrid_cache": {
    "enabled": false
  },
  "mirror": {
    "enabled": true
  }
},
]
"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."
    }
  }
]
}

```

```

],
"_links": {
  "self": {
    "href": "/api/storage/aggregates?recommend=true&fields=*"
  }
}
}
}

```

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 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 |
|-----------|---------|-------|----------|---|
| 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'. |

| Name | Type | In | Required | Description |
|---------------------------|---------|-------|----------|---|
| 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'. |
| node.name | string | query | False | Filter by node.name |
| node.uuid | string | query | False | Filter by node.uuid |
| home_node.name | string | query | False | Filter by home_node.name |
| home_node.uuid | string | query | False | Filter by home_node.uuid |
| name | string | query | False | Filter by name |
| statistics.iops_raw.total | integer | query | False | Filter by statistics.iops_raw.total • Introduced in: 9.7 |
| statistics.iops_raw.other | integer | query | False | Filter by statistics.iops_raw.other • Introduced in: 9.7 |
| statistics.iops_raw.read | integer | query | False | Filter by statistics.iops_raw.read • Introduced in: 9.7 |
| statistics.iops_raw.write | integer | query | False | Filter by statistics.iops_raw.write • Introduced in: 9.7 |

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

| Name | Type | In | Required | Description |
|---|---------|-------|----------|---|
| statistics.latency_read.other | integer | query | False | Filter by statistics.latency_read.other • Introduced in: 9.7 |
| statistics.latency_read.read | integer | query | False | Filter by statistics.latency_read.read • Introduced in: 9.7 |
| statistics.latency_read.write | integer | query | False | Filter by statistics.latency_read.write • 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 |
| 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.available | integer | query | False | Filter by space.block_storage.available |
| space.block_storage.inactive_user_data | integer | query | False | Filter by space.block_storage.inactive_user_data |

| Name | Type | In | Required | Description |
|---|---------|-------|----------|---|
| space.block_storage.used | integer | query | False | Filter by space.block_storage.used |
| space.block_storage.physical_used | integer | query | False | Filter by space.block_storage.physical_used • Introduced in: 9.9 |
| space.cloud_storage.used | integer | query | False | Filter by space.cloud_storage.used |
| space.footprint | integer | query | False | Filter by space.footprint |
| space.encyency_without_snapshots.savings | integer | query | False | Filter by space.encyency_without_snapshots.savings |
| space.encyency_without_snapshots.ratio | number | query | False | Filter by space.encyency_without_snapshots.ratio |
| space.encyency_without_snapshots.logical_used | integer | query | False | Filter by space.encyency_without_snapshots.logical_used |
| space.encyency.savings | integer | query | False | Filter by space.encyency.savings |
| space.encyency.ratio | number | query | False | Filter by space.encyency.ratio |
| space.encyency.logical_used | integer | query | False | Filter by space.encyency.logical_used |

| Name | Type | In | Required | Description |
|--|---------|-------|----------|--|
| space.encyency_wit hout_snapshots_flex clones.savings | integer | query | False | Filter by space.encyency_wit hout_snapshots_flex clones.savings • Introduced in: 9.9 |
| space.encyency_wit hout_snapshots_flex clones.ratio | number | query | False | Filter by space.encyency_wit hout_snapshots_flex clones.ratio • Introduced in: 9.9 |
| space.encyency_wit hout_snapshots_flex clones.logical_used | integer | query | False | Filter by space.encyency_wit hout_snapshots_flex clones.logical_used • Introduced in: 9.9 |
| inactive_data_reporti ng.enabled | boolean | query | False | Filter by inactive_data_reporti ng.enabled • Introduced in: 9.8 |
| inactive_data_reporti ng.start_time | string | query | False | Filter by inactive_data_reporti ng.start_time • Introduced in: 9.8 |
| uuid | string | query | False | Filter by uuid |
| state | string | query | False | Filter by state |
| block_storage.hybrid _cache.used | integer | query | False | Filter by block_storage.hybrid _cache.used |
| block_storage.hybrid _cache.enabled | boolean | query | False | Filter by block_storage.hybrid _cache.enabled |

| Name | Type | In | Required | Description |
|---------------------------------------|---------|-------|----------|---|
| block_storage.hybrid_cache.raid_type | string | query | False | Filter by block_storage.hybrid_cache.raid_type |
| block_storage.hybrid_cache.disk_count | integer | query | False | Filter by block_storage.hybrid_cache.disk_count |
| block_storage.hybrid_cache.size | integer | query | False | Filter by block_storage.hybrid_cache.size |
| block_storage.mirror.enabled | boolean | query | False | Filter by block_storage.mirror.enabled |
| block_storage.mirror.state | string | query | False | Filter by block_storage.mirror.state |
| block_storage.primary.disk_class | string | query | False | Filter by block_storage.primary.disk_class |
| block_storage.primary.checksum_style | string | query | False | Filter by block_storage.primary.checksum_style |
| block_storage.primary.disk_count | integer | query | False | Filter by block_storage.primary.disk_count |
| block_storage.primary.disk_type | string | query | False | Filter by block_storage.primary.disk_type • Introduced in: 9.7 |
| 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 |

| Name | Type | In | Required | Description |
|---------------------------|---------|-------|----------|---|
| block_storage.plexes.name | string | query | False | Filter by block_storage.plexes.name |
| dr_home_node.uuid | string | query | False | Filter by dr_home_node.uuid |
| dr_home_node.name | string | query | False | Filter by dr_home_node.name |
| snaplock_type | string | query | False | Filter by snaplock_type |
| create_time | string | query | False | Filter by create_time |
| metric.timestamp | string | query | False | Filter by metric.timestamp • Introduced in: 9.7 |
| metric.throughput.total | integer | query | False | Filter by metric.throughput.total • Introduced in: 9.7 |
| metric.throughput.other | integer | query | False | Filter by metric.throughput.other • Introduced in: 9.7 |
| metric.throughput.read | integer | query | False | Filter by metric.throughput.read • 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.latency.total | integer | query | False | Filter by metric.latency.total • Introduced in: 9.7 |
| metric.latency.other | integer | query | False | Filter by metric.latency.other • Introduced in: 9.7 |
| metric.latency.read | integer | query | False | Filter by metric.latency.read • Introduced in: 9.7 |
| metric.latency.write | integer | query | False | Filter by metric.latency.write • Introduced in: 9.7 |
| metric.iops.total | integer | query | False | Filter by metric.iops.total • Introduced in: 9.7 |
| metric.iops.other | integer | query | False | Filter by metric.iops.other • 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.write | integer | query | False | Filter by metric.iops.write • Introduced in: 9.7 |
| metric.duration | string | query | False | Filter by metric.duration • Introduced in: 9.7 |
| metric.status | string | query | False | Filter by metric.status • 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 | |
| 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"
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "block_storage": {
    "hybrid_cache": {
      "disk_count": "6",
      "raid_type": "raid_dp",
      "size": "1612709888",
      "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"
  }
},
"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": "non_snaplock",
  "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": "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": "2017-01-25T11:20:13Z"
},
"uuid": "string"
},
"spares": {
  "checksum_style": "block",
  "disk_class": "solid_state",
  "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",
"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. |
| 19726344 | No recommendation can be made for this cluster. |
| 19726357 | Aggregate recommendations are not supported on MetroCluster with Fibre Channel (FC). |

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

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. |

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

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

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

| 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 | |
| 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 |

raid_group

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

| Name | Type | Description |
|---------------------|---|--|
| disk_class | string | Disk class of spares |
| layout_requirements | array[layout_requirement] | Available RAID protections and their restrictions |
| node | node | Node where the spares are assigned |
| size | integer | Usable size of each spare in bytes |
| syncmirror_pool | string | SyncMirror spare pool |
| 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 |

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. |

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

| Name | Type | In | Required | Description |
|----------------|---------|-------|----------|---|
| 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 | |
| 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. |

| Name | Type | Description |
|---------------|----------------------------|---|
| 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": "raid_dp",
      "size": "1612709888",
      "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"
    }
  },
  "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": "non_snaplock",
"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": "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": "2017-01-25T11:20:13Z"
  },
  "uuid": "string"
}

```

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 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. |
| 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. |

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

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. |

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.