Overview

REST API reference

Table of Contents

Overview

ONTAP REST API Storage aggregates uuid endpoints

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,
    "physical_used_percent": 1,
    "volume_footprints_percent": 14,
    "aggregate_metadata": 2655,
    "aggregate_metadata_percent": 8,
    "used_including_snapshot_reserve": 674685,
    "used_including_snapshot_reserve_percent": 35,
    "data_compacted_count": 666666,
    "data_compaction_space_saved": 654566,
    "data_compaction_space_saved_percent": 47,
    "volume_deduplication_shared_count": 567543,
    "volume_deduplication_space_saved": 23765,
    "volume_deduplication_space_saved_percent": 32
  },
  "snapshot": {
```

```
      "used_percent": 45,
      "available": 2000,
      "total": 5000,
      "used": 3000,
      "reserve_percent": 20
    },
    "cloud_storage": {
      "used": 0
    },
    "efficiency": {
      "savings": 1408029,
      "ratio": 6.908119720880661,
      "logical_used": 1646350
    },
    "efficiency_without_snapshots": {
      "savings": 0,
      "ratio": 1,
      "logical_used": 737280
    },
    "efficiency_without_snapshots_flexclones": {
      "savings": 5000,
      "ratio": 2,
      "logical_used": 10000
    }
  },
  "snapshot": {
    "files_total": 10,
    "files_used": 3,
    "max_files_available": 5,
    "max_files_used": 50
  },
  "state": "online",
  "snaplock_type": "non_snaplock",
  "create_time": "2018-12-04T15:40:38-05:00",
  "data_encryption": {
    "software_encryption_enabled": false,
    "drive_protection_enabled": false
  },
  "block_storage": {
    "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-
9590-5a402bffbe4b?fields=data_encryption.software_encryption_enabled" -H
"accept: application/json"

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

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

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

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

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

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

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

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

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

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`

DELETE `/storage/aggregates/{uuid}`

Introduced In: 9.6

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

Related ONTAP commands

* `storage aggregate delete`

Parameters

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

Response

```
Status: 202, Accepted
```

Name	Type	Description
job	job_link	

Example response

```json
{
  "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.

Error Code	Description
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.