



Manage specific storage pools

REST API reference

NetApp
September 09, 2025

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

Table of Contents

- Manage specific storage pools 1
 - Storage pools UUID endpoint overview 1
 - Updating storage pools 1
 - Deleting storage pools 1
 - Examples 1
 - Delete a storage pool specified by the UUID 10
 - Related ONTAP commands 11
 - Parameters 11
 - Response 11
 - Response 12
 - Error 12
 - Definitions 13
 - Retrieve a storage pool specified by the UUID 14
 - Related ONTAP commands 15
 - Parameters 15
 - Response 15
 - Error 18
 - Definitions 18
 - Update a storage pool specified by the UUID 22
 - Related ONTAP commands 23
 - Parameters 23
 - Request Body 24
 - Response 26
 - Response 28
 - Error 28
 - Definitions 29

Manage specific storage pools

Storage pools UUID endpoint overview

Updating storage pools

The PATCH operation is used to modify properties of the storage pool. There are several properties that can be modified on a storage pool. PATCH operations on a storage pool are restricted when another PATCH operation is in progress. The following is a list of properties that can be modified using the PATCH operation including a brief description for each:

- name - Can be updated to rename the storage pool.
- capacity.disk_count - Can be updated to increase the number of disks in a storage pool.
- capacity.spare_allocation_units[].count - Modifying this value requires that the user specify capacity.spare_allocation_units[].node as well. Modifying this value redistributes spare cache capacity among the nodes specified in the operation. When expanding a storage pool, the cache tiers of all aggregates using the storage pool's allocation units are expanded automatically.

Simulated storage pool expansion

The PATCH operation also supports simulated expansion of a storage pool. Running PATCH with the query "simulate" set to "true", and "capacity.disk_count" set to the final disk count will return a response containing the projected new capacity and the new constituent disk list for the storage pool.

Deleting storage pools

If cache capacity from a storage pool is being used in an aggregate, it cannot be deleted. See the /storage/aggregates API for details on deleting aggregates.

Examples

Retrieving a specific pool from the cluster

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

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

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

# The response:
{
  "uuid": "8255fef7-4737-11ec-bd1b-005056bbb879",
  "name": "new_sp",
```

```

"storage_type": "ssd",
"nodes": [
  {
    "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
    "name": "node-1",
  },
  {
    "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
    "name": "node-2",
  }
],
"capacity": {
  "remaining": 1846542336,
  "total": 7386169344,
  "spare_allocation_units": [
    {
      "node": {
        "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
        "name": "node-1",
      },
      "count": 1,
      "syncmirror_pool": "pool0",
      "size": 1846542336,
      "available_size": 1846542336
    },
    {
      "node": {
        "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
        "name": "node-2",
      },
      "count": 0,
      "syncmirror_pool": "pool0",
      "size": 1846542336,
      "available_size": 0
    }
  ],
  "used_allocation_units": [
    {
      "aggregate": {
        "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
        "name": "test_a"
      },
      "allocated_unit_count": 2,
      "node": {
        "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
        "name": "node-1",
      }
    }
  ]
}

```

```

    },
    "capacity": 2769813504
  },
  {
    "aggregate": {
      "uuid": "f4cc30d5-b052-493a-a49f-19781425f987",
      "name": "test_b"
    },
    "allocated_unit_count": 1,
    "node": {
      "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
      "name": "node-2",
    },
    "capacity": 1384906752
  }
],
"disk_count": 4,
"disks": [
  {
    "disk": {
      "name": "VMw-1.11"
    },
    "usable_size": 1902379008,
    "total_size": 1908871168,
  },
  {
    "disk": {
      "name": "VMw-1.12"
    },
    "usable_size": 1902379008,
    "total_size": 1908871168,
  },
  {
    "disk": {
      "name": "VMw-1.23"
    },
    "usable_size": 1902379008,
    "total_size": 1908871168,
  },
  {
    "disk": {
      "name": "VMw-1.24"
    },
    "usable_size": 1902379008,
    "total_size": 1908871168,
  }
]

```

```

    ]
  },
  "health": {
    "state": "normal",
    "is_healthy": true
  },
}

```

Simulating storage pool expansion

The following example shows the response for a simulated storage pool expansion based on the values of the 'capacity.disk_count' attribute passed in. The query does not modify the existing storage pool, but rather returns how it will look after the expansion. This will be reflected in the following attributes:

- capacity.total- Total space, in bytes.
- capacity.remaining - New remaining capacity, in bytes.
- capacity.disks.disk - New list of constituent disks.
- capacity.disk_count - New number of disks in the pool.

```

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/pools/cae60cfe-deae-42bd-babb-ef437d118314?simulate=true" -H "accept: application/json" -d
{"\"capacity\": {\"disk_count\": 6}}"

# The response:
{
  "records": [
    {
      "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
      "name": "new_sp",
      "capacity": {
        "remaining": 1846542336,
        "total": 7386169344,
        "used_allocation_units": [
          {
            "aggregate": {
              "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
              "name": "test_a"
            },
            "current_usage": 2769813504
          },
          {
            "aggregate": {

```

```

        "uuid": "f4cc30d5-b052-493a-a49f-19781425f987",
        "name": "test_b"
    },
    "current_usage": 1384906752
}
],
"disk_count": 6,
"disks": [
    {
        "disk": {
            "name": "VMw-1.11"
        },
    },
    {
        "disk": {
            "name": "VMw-1.12"
        },
    },
    {
        "disk": {
            "name": "VMw-1.23"
        },
    },
    {
        "disk": {
            "name": "VMw-1.24"
        },
    },
    {
        "disk": {
            "name": "VMw-1.13"
        },
    },
    {
        "disk": {
            "name": "VMw-1.14"
        },
    },
]
}
]
}

```

Adding capacity to a storage pool

The following example shows the workflow of adding disks to the storage pool.

Step 1: Check the current disk count on the storage pool.

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

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

# The response:
{
  "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
  "name": "sp1",
  "capacity": {
    "disk_count": 4
  }
}
```

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

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/pools/19425837-f2fa-4a9f-8f01-712f626c983c" -H "accept: application/hal+json" -d '{"capacity":{"disk_count": 6}}'

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

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

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

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

# The response:
{
  "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
  "name": "sp1",
  "capacity": {
    "disk_count": 6
  }
}
```

The following example shows the workflow to redistribute spare capacity among nodes sharing the storage pool Step 1: Check the current spare capacity distribution of the pool.

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

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

# The response:
{
  "uuid": "f3aafdc6-be35-4d93-9590-5a402bffb4b",
  "name": "sp1",
  "capacity": {
    "spare_allocation_units": [
      {
        "node": {
          "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
          "name": "node-1",
        },
        "count": 1,
        "syncmirror_pool": "pool0",
        "size": 1846542336,
        "available_size": 1846542336
      },
      {
        "node": {
          "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
          "name": "node-2",
        },
        "count": 0,
        "syncmirror_pool": "pool0",
        "size": 1846542336,
        "available_size": 0
      }
    ],
  }
}
```

Step 2: Update the pool so that the spare allocation unit count is symmetrically modified for each node. The response to PATCH is a job unless the request is invalid.

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

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/pools/f3aafdc6-be35-4d93-9590-5a402bffb4b" -H "accept: application/hal+json" -d '{ "capacity" : { "spare_allocation_units": [{"node": {"name": "node-1"}, "count": 0}, {"node": {"name": "node-2"}, "count": 1}]}}'

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

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

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

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

# The response:
{
  "uuid": "f3aafdc6-be35-4d93-9590-5a402bffb4b",
  "name": "sp1",
  "capacity": {
    "spare_allocation_units": [
      {
        "node": {
          "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
          "name": "node-1",
        },
        "count": 0,
        "syncmirror_pool": "pool0",
        "size": 1846542336,
        "available_size": 0
      },
      {
        "node": {
          "uuid": "cf9ab500-ff3e-4bce-bfd7-d679e6078f47",
          "name": "node-2",
        },
        "count": 1,
        "syncmirror_pool": "pool0",
        "size": 1846542336,
        "available_size": 1846542336
      }
    ]
  }
}
```

Delete a storage pool specified by the UUID

DELETE /storage/pools/{uuid}

Introduced In: 9.11

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

Related ONTAP commands

- `storage pool delete`

Parameters

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

Response

Status: 200, Ok

Name	Type	Description
job	job_link	

Example response

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

Response

Status: 202, Accepted

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
11209668	Capacity from pool is being used by one or more aggregates.
11209670	Unable to stop sharing a constituent disk of the storage pool.
11209671	Cannot find a node sharing this storage pool.
11209672	Disk is not shared.
11209673	Unable to assign shared capacity from a constituent disk.
11209675	A constituent disk of the storage pool has non-spare capacity.
11209676	A constituent disk of the storage pool is reserved for coredump.

Also see the table of common errors in the [Response body](#) overview section of this documentation.

Name	Type	Description
error	returned_error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

returned_error

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

Retrieve a storage pool specified by the UUID

GET /storage/pools/{uuid}

Introduced In: 9.11

Retrieves the storage pool specified by the UUID.

Related ONTAP commands

- `storage pool show -storage-pool-uuid`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Storage pool UUID.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
capacity	capacity	
health	health	Properties that outline shared storage pool health.
name	string	Storage pool name.
nodes	array[node_reference]	Nodes that can use this storage pool for their aggregates.
storage_type	string	Storage type for the disks used to create the storage pool.
uuid	string	Storage pool UUID.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "capacity": {
    "disks": [
      {
        "disk": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "1.0.1"
        },
        "total_size": 0,
        "usable_size": 0
      }
    ],
    "remaining": 0,
    "spare_allocation_units": [
      {
        "available_size": 0,
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "size": 0,
        "syncmirror_pool": "string"
      }
    ],
    "total": 0,
    "used_allocation_units": [
      {
        "aggregate": {
          "_links": {
            "self": {
```

```

        "href": "/api/resourcelink"
    },
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"current_usage": 0,
"node": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
}
]
},
"health": {
    "state": "string",
    "unhealthy_reason": {
        "arguments": [
            {
                "code": "string",
                "message": "string"
            }
        ],
        "code": "4",
        "message": "entry doesn't exist"
    }
},
"name": "string",
"nodes": [
    {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
],
"storage_type": "string",
"uuid": "string"

```

```
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
11206662	There is no storage pool matching the specified UUID or name.
11215856	The specified storage pool is not healthy.

Also see the table of common errors in the [Response body](#) overview section of this documentation.

Name	Type	Description
error	returned_error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

disk

Reference to the constituent disk object.

Name	Type	Description
_links	_links	
name	string	

storage_pool_disk

Name	Type	Description
disk	disk	Reference to the constituent disk object.
total_size	integer	Raw capacity of the disk, in bytes.
usable_size	integer	Usable capacity of this disk, in bytes.

node

Specifies what node can use this set of allocation units.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pool_spare_allocation_unit

Name	Type	Description
available_size	integer	The usable capacity of this set of allocation units.
count	integer	The number of spare allocation units on this node.
node	node	Specifies what node can use this set of allocation units.
size	integer	Size of each allocation unit.
syncmirror_pool	string	The RAID SyncMirror Pool to which this allocation unit is assigned.

aggregate

The aggregate that is using this cache capacity.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

node

The node hosting the aggregate using this set of allocation units.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_pool_used_allocation_unit

Name	Type	Description
aggregate	aggregate	The aggregate that is using this cache capacity.
count	integer	The number of allocation units used by this aggregate.

Name	Type	Description
current_usage	integer	The amount of cache space used by this aggregate.
node	node	The node hosting the aggregate using this set of allocation units.

capacity

Name	Type	Description
disk_count	integer	The number of disks in the storage pool.
disks	array[storage_pool_disk]	Properties of each disk used in the shared storage pool.
remaining	integer	Remaining usable capacity in the flash pool, in bytes.
spare_allocation_units	array[storage_pool_spare_allocation_unit]	Properties of spare allocation units.
total	integer	Total size of the flash pool, in bytes.
used_allocation_units	array[storage_pool_used_allocation_unit]	Information about the storage pool allocation units participating in the cache tier of an aggregate.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Indicates why the storage pool is unhealthy. This property is not returned for healthy storage pools.

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

Name	Type	Description
message	string	Error message

health

Properties that outline shared storage pool health.

Name	Type	Description
is_healthy	boolean	Indicates whether the storage pool is able to participate in provisioning operations.
state	string	The state of the shared storage pool.
unhealthy_reason	error	Indicates why the storage pool is unhealthy. This property is not returned for healthy storage pools.

node_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

returned_error

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

Update a storage pool specified by the UUID

PATCH /storage/pools/{uuid}

Introduced In: 9.11

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

Related ONTAP commands

- `storage pool rename`
- `storage pool reassign`
- `storage pool add`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Storage pool UUID.
simulate	boolean	query	False	When set to "true" in conjunction with a PATCH on "capacity.disk_count", the end point returns a simulated layout of the storage pool with the additional disks, without changing system state. • Introduced in: 9.12

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

Request Body

Name	Type	Description
capacity	capacity	
name	string	Storage pool name.
nodes	array[node_reference]	Nodes that can use this storage pool for their aggregates.
storage_type	string	Storage type for the disks used to create the storage pool.
uuid	string	Storage pool UUID.

Example request

```
{
  "capacity": {
    "disks": [
      {
        "disk": {
          "name": "1.0.1"
        },
        "total_size": 0,
        "usable_size": 0
      }
    ],
    "remaining": 0,
    "spare_allocation_units": [
      {
        "available_size": 0,
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "size": 0,
        "syncmirror_pool": "string"
      }
    ],
    "total": 0,
    "used_allocation_units": [
      {
        "aggregate": {
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "current_usage": 0,
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ]
  },
  "name": "string",
  "nodes": [
    {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

```
],  
  "storage_type": "string",  
  "uuid": "string"  
}
```

Response

Status: 200, Ok

Name	Type	Description
job	job_link	
num_records	integer	Number of shared storage pools in the cluster.
records	array[storage_pool]	

Example response

```
{
  "job": {
    "uuid": "string"
  },
  "num_records": 1,
  "records": [
    {
      "capacity": {
        "disks": [
          {
            "disk": {
              "name": "1.0.1"
            },
            "total_size": 0,
            "usable_size": 0
          }
        ],
        "remaining": 0,
        "spare_allocation_units": [
          {
            "available_size": 0,
            "node": {
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "size": 0,
            "syncmirror_pool": "string"
          }
        ],
        "total": 0,
        "used_allocation_units": [
          {
            "aggregate": {
              "name": "aggr1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "current_usage": 0,
            "node": {
              "name": "node1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            }
          }
        ]
      }
    }
  ],
}
```

```

    "name": "string",
    "nodes": [
      {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "storage_type": "string",
    "uuid": "string"
  }
]
}

```

Response

Status: 202, Accepted

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
11211658	Node does not have enough spare capacity.
11211659	Valid allocation unit input is required.
11211662	Specified node is not part of the storage pool.
11211663	Failed to reassign available capacity in the storage pool.
11211664	Could not fix the broken allocation unit for the storage pool.
11212673	Could not grow one or more aggregates.
11212679	Adding specified number of disks will expand storage pool beyond maximum supported disk limit.
11212680	Incorrect node specified.
11212681	0 is an invalid value for disk_count.

Error Code	Description
11212682	Adding the specified number of disks will result in the storage pool reaching the maximum disk limit reserved for RAID-TEC use only. At this limit, the storage pool can only allocate capacity to aggregates containing RAID-TEC RAID groups. Existing aggregates containing RAID groups other than RAID-TEC will not automatically grow to the new capacity.
11212683	Renaming storage pool to new name failed.
11212763	Storage pool add job failed.
11215657	Storage pool PATCH request have missing parameters.
11215658	Storage pool PATCH request for reassign is invalid.
11215659	Storage pool PATCH request for reassign have invalid allocation unit count.
11215660	Storage pool PATCH request for reassign have invalid node name.
11215662	Storage pool PATCH request have invalid disk count.

Also see the table of common errors in the [Response body](#) overview section of this documentation.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

disk

Reference to the constituent disk object.

Name	Type	Description
name	string	

storage_pool_disk

Name	Type	Description
disk	disk	Reference to the constituent disk object.
total_size	integer	Raw capacity of the disk, in bytes.
usable_size	integer	Usable capacity of this disk, in bytes.

node

Specifies what node can use this set of allocation units.

Name	Type	Description
name	string	
uuid	string	

storage_pool_spare_allocation_unit

Name	Type	Description
available_size	integer	The usable capacity of this set of allocation units.
count	integer	The number of spare allocation units on this node.

Name	Type	Description
node	node	Specifies what node can use this set of allocation units.
size	integer	Size of each allocation unit.
syncmirror_pool	string	The RAID SyncMirror Pool to which this allocation unit is assigned.

aggregate

The aggregate that is using this cache capacity.

Name	Type	Description
name	string	
uuid	string	

node

The node hosting the aggregate using this set of allocation units.

Name	Type	Description
name	string	
uuid	string	

storage_pool_used_allocation_unit

Name	Type	Description
aggregate	aggregate	The aggregate that is using this cache capacity.
count	integer	The number of allocation units used by this aggregate.
current_usage	integer	The amount of cache space used by this aggregate.
node	node	The node hosting the aggregate using this set of allocation units.

capacity

Name	Type	Description
disk_count	integer	The number of disks in the storage pool.
disks	array[storage_pool_disk]	Properties of each disk used in the shared storage pool.
remaining	integer	Remaining usable capacity in the flash pool, in bytes.
spare_allocation_units	array[storage_pool_spare_allocation_unit]	Properties of spare allocation units.
total	integer	Total size of the flash pool, in bytes.
used_allocation_units	array[storage_pool_used_allocation_unit]	Information about the storage pool allocation units participating in the cache tier of an aggregate.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Indicates why the storage pool is unhealthy. This property is not returned for healthy storage pools.

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

health

Properties that outline shared storage pool health.

Name	Type	Description
is_healthy	boolean	Indicates whether the storage pool is able to participate in provisioning operations.
state	string	The state of the shared storage pool.

node_reference

Name	Type	Description
name	string	
uuid	string	

storage_pool

Name	Type	Description
capacity	capacity	
name	string	Storage pool name.
nodes	array[node_reference]	Nodes that can use this storage pool for their aggregates.
storage_type	string	Storage type for the disks used to create the storage pool.
uuid	string	Storage pool UUID.

job_link

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

returned_error

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

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

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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