



Manage QoS policies

ONTAP 9.14.1 REST API reference

NetApp
May 23, 2024

Table of Contents

- Manage QoS policies 1
 - Storage Qos policies endpoint overview 1
 - Retrieve QoS policies 2
 - Create a QoS policy 13
 - Delete a QoS policy 22
 - Retrieve a QoS policy 25
 - Update a QoS policy 30

Manage QoS policies

Storage QoS policies endpoint overview

Quality of Service Configuration

A QoS policy defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy is associated. There are two types of policies that can be configured: fixed, which defines a fixed SLO, or adaptive which defines a variable SLO for a storage object. Adaptive policies vary the SLO depending on the space usage of the storage object. A policy can be either a fixed policy or an adaptive one, not both. Service level objectives include minimum and maximum limits on throughput in terms of IOPS. Only maximum limits can be set in terms of both IOPS and/or throughput (MB/s). A QoS policy can be used to enforce SLOs for multiple storage objects by specifying "capacity_shared" to true. For example, if a QoS policy with "capacity_shared" is set to true and it has maximum_throughput_iops set to 1000, and this policy is assigned to four volumes, then the combined throughput of all four volumes is limited to 1000 IOPS. If "capacity_shared" is set to false then, each storage object will have its SLOs enforced individually. For example, in the previous case if the same policy was applied to four volumes but with "capacity_shared" set to false, then each of the volumes would be limited to 1000 IOPS individually. Once "capacity_shared" is set, it cannot be modified. Adaptive parameters can specify the variable SLOs in terms of IOPS/TB. The actual IOPS enforced on the storage object can be calculated using the allocated space on the storage object. The policies are enforced individually amongst storage objects.

Examples

1) Create a fixed QoS policy

The following example shows how to create a fixed QoS policy to limit throughput for a storage object between 5000 IOPS and 10000 IOPS which has capacity_shared set to false. This QoS policy can be used as a template to apply on multiple storage objects to provide individual SLOs to each object.

```
curl -X POST
"https://172.21.69.245/api/storage/qos/policies?return_timeout=0" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
  \"fixed\": { \"capacity_shared\": false, \"max_throughput_iops\": 10000,
  \"min_throughput_iops\": 5000 }, \"name\":
  \"qos_policy_5000_to_10000_iops\", \"svm\": { \"name\": \"vs0\" }}"
```

2) Create an adaptive QoS policy

The following example shows how to create an adaptive QoS policy which provides 5000 IOPS per GB of allocated space for a storage object with a peak of 6000 IOPS. Minimum IOPS regardless of allocated space are 1000 IOPS.

```
curl -X POST
"https://172.21.69.245/api/storage/qos/policies?return_timeout=0" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"adaptive\": { \"absolute_min_iops\": 1000, \"expected_iops\": 5000,
\"expected_iops_allocation\": \"used_space\", \"peak_iops\": 6000,
\"peak_iops_allocation\": \"allocated_space\" }, \"name\":
\"adaptive_pg_5k_to_6k\", \"svm\": { \"name\": \"vs0\" }}"
```

3) Update an existing QoS policy

The following example shows how to update SLOs of an existing QoS policy and also rename it.

```
curl -X PATCH "https://172.21.69.245/api/storage/qos/policies/d38bafc0-
5a51-11e9-bd5b-005056ac6f1f?return_timeout=0" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"fixed\": {
\"max_throughput_iops\": 15000, \"min_throughput_iops\": 10000 },
\"name\": \"qos_policy_10k_to_15k_iops\"}"
```

4) Delete an existing QoS policy

When a QoS policy is deleted any associations of the policy with a storage objects are also removed.

```
curl -X DELETE "https://172.21.69.245/api/storage/qos/policies/d38bafc0-
5a51-11e9-bd5b-005056ac6f1f?return_timeout=0" -H "accept:
application/json"
```

Retrieve QoS policies

GET /storage/qos/policies

Introduced In: 9.6

Retrieves a collection of QoS policies.

Parameters

| Name | Type | In | Required | Description |
|---------------------------|---------|-------|----------|--|
| pgid | integer | query | False | Filter by pgid <ul style="list-style-type: none"> • Introduced in: 9.10 |
| fixed.max_throughput_mbps | integer | query | False | Filter by fixed.max_throughput_mbps <ul style="list-style-type: none"> • Max value: 4194303 • Min value: 0 |
| fixed.min_throughput_mbps | integer | query | False | Filter by fixed.min_throughput_mbps <ul style="list-style-type: none"> • Introduced in: 9.8 • Max value: 4194303 • Min value: 0 |
| fixed.capacity_shard | boolean | query | False | Filter by fixed.capacity_shard |
| fixed.min_throughput_iops | integer | query | False | Filter by fixed.min_throughput_iops <ul style="list-style-type: none"> • Max value: 2147483647 • Min value: 0 |
| fixed.max_throughput_iops | integer | query | False | Filter by fixed.max_throughput_iops <ul style="list-style-type: none"> • Max value: 2147483647 • Min value: 0 |

| Name | Type | In | Required | Description |
|-----------------------------------|---------|-------|----------|--|
| policy_class | string | query | False | Filter by policy_class • Introduced in: 9.10 |
| name | string | query | False | Filter by name |
| scope | string | query | False | Filter by scope • Introduced in: 9.11 |
| uuid | string | query | False | Filter by uuid |
| object_count | integer | query | False | Filter by object_count |
| svm.uuid | string | query | False | Filter by svm.uuid |
| svm.name | string | query | False | Filter by svm.name |
| adaptive.block_size | string | query | False | Filter by adaptive.block_size • Introduced in: 9.10 |
| adaptive.expected_iops_allocation | string | query | False | Filter by adaptive.expected_iops_allocation • Introduced in: 9.10 |
| adaptive.peak_iops_allocation | string | query | False | Filter by adaptive.peak_iops_allocation • Introduced in: 9.10 |
| adaptive.expected_iops | integer | query | False | Filter by adaptive.expected_iops |
| adaptive.peak_iops | integer | query | False | Filter by adaptive.peak_iops |

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

Response

Status: 200, Ok

| Name | Type | Description |
|-------------|-------------------------------------|-------------------|
| _links | _links | |
| error | error | |
| num_records | integer | Number of records |
| records | array[qos_policy] | |

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist"
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "adaptive": {
      "block_size": "any",
      "expected_iops_allocation": "used_space",
      "peak_iops_allocation": "used_space"
    },
    "name": "extreme",
    "object_count": 0,
    "pgid": 0,
    "policy_class": "undefined",
    "scope": "cluster",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

```
}  
}
```

Error

Status: Default, Error

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

_links

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

adaptive

Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

| Name | Type | Description |
|--------------------------|---------|---|
| absolute_min_iops | integer | Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set. |
| block_size | string | Specifies the block size |
| expected_iops | integer | Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set. |
| expected_iops_allocation | string | Specifies the size to be used to calculate expected IOPS per TB. The size options are either the storage object allocated space or the storage object used space. |
| peak_iops | integer | Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size. |
| peak_iops_allocation | string | Specifies the size to be used to calculate peak IOPS per TB. The size options are either the storage object allocated space or the storage object used space. |

fixed

QoS policy-groups define a fixed service level objective (SLO) for a storage object.

| Name | Type | Description |
|-----------------|---------|--|
| capacity_shared | boolean | Specifies whether the capacities are shared across all objects that use this QoS policy-group. Default is false. |

| Name | Type | Description |
|---------------------|---------|---|
| max_throughput_iops | integer | Maximum throughput defined by this policy. It is specified in terms of IOPS. 0 means no maximum throughput is enforced. |
| max_throughput_mbps | integer | Maximum throughput defined by this policy. It is specified in terms of Mbps. 0 means no maximum throughput is enforced. |
| min_throughput_iops | integer | Minimum throughput defined by this policy. It is specified in terms of IOPS. 0 means no minimum throughput is enforced. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set. |
| min_throughput_mbps | integer | Minimum throughput defined by this policy. It is specified in terms of Mbps. 0 means no minimum throughput is enforced. |

svm

SVM, applies only to SVM-scoped objects.

| Name | Type | Description |
|--------|------------------------|---|
| _links | _links | |
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

qos_policy

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

| Name | Type | Description |
|--------------|--------------------------|---|
| adaptive | adaptive | Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space. |
| fixed | fixed | QoS policy-groups define a fixed service level objective (SLO) for a storage object. |
| name | string | Name of the QoS policy. |
| object_count | integer | Number of objects attached to this policy. |
| pgid | integer | Policy group ID of the QoS policy. |
| policy_class | string | Class of the QoS policy. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| 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. |

Create a QoS policy

POST /storage/qos/policies

Introduced In: 9.6

Creates a QoS policy.

Required properties

- `svm.uuid` or `svm.name` - The existing SVM owning the QoS policy.
- `name` - The name of the QoS policy.
- `fixed.*` or `adaptive.*` - Either of the fixed or adaptive parameters.

Default property values

- If `fixed.*` parameters are specified, then `capacity.shared` is set to `false` by default.

Related ONTAP commands

- `qos policy-group create`
- `qos adaptive-policy-group create`

Parameters

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

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

Request Body

| Name | Type | Description |
|----------|--------------------------|---|
| _links | _links | |
| adaptive | adaptive | Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space. |
| fixed | fixed | QoS policy-groups define a fixed service level objective (SLO) for a storage object. |

| Name | Type | Description |
|--------------|---------|---|
| name | string | Name of the QoS policy. |
| object_count | integer | Number of objects attached to this policy. |
| pgid | integer | Policy group ID of the QoS policy. |
| policy_class | string | Class of the QoS policy. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | |

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "adaptive": {
    "block_size": "any",
    "expected_iops_allocation": "used_space",
    "peak_iops_allocation": "used_space"
  },
  "name": "extreme",
  "object_count": 0,
  "pgid": 0,
  "policy_class": "undefined",
  "scope": "cluster",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 202, Accepted

| Name | Type | Description |
|------|--------------------------|-------------|
| job | job_link | |

Example response

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

Headers

| Name | Description | Type |
|----------|---|--------|
| Location | Useful for tracking the resource location | string |

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

| Error Code | Description |
|------------|---|
| 8454147 | The maximum limit for QoS policies has been reached. |
| 8454154 | The name specified for creating conflicts with an existing QoS policy name. |
| 8454194 | The minimum throughput value for the policy group must be less than or equal to the maximum throughput value. |

| Error Code | Description |
|------------|--|
| 8454260 | Invalid value for maximum and minimum fields. Valid values for max_throughput_iops and max_throughput_mbps combination is for the ratio of max_throughput_mbps and max_throughput_iops to be within 1 to 4096. |
| 8454273 | Invalid value for an adaptive field. Value should be non-zero. |
| 8454274 | Invalid value for an adaptive field. Value for expected_iops must be less than or equal to the value for peak_iops. |
| 8454277 | The name specified for creating an adaptive QoS policy conflicts with an existing fixed QoS policy name. |
| 8454278 | The name specified for creating a fixed QoS policy conflicts with an existing adaptive QoS policy name. |

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 | |

adaptive

Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

| Name | Type | Description |
|--------------------------|---------|---|
| absolute_min_iops | integer | Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set. |
| block_size | string | Specifies the block size |
| expected_iops | integer | Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set. |
| expected_iops_allocation | string | Specifies the size to be used to calculate expected IOPS per TB. The size options are either the storage object allocated space or the storage object used space. |
| peak_iops | integer | Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size. |

| Name | Type | Description |
|----------------------|--------|---|
| peak_iops_allocation | string | Specifies the size to be used to calculate peak IOPS per TB. The size options are either the storage object allocated space or the storage object used space. |

fixed

QoS policy-groups define a fixed service level objective (SLO) for a storage object.

| Name | Type | Description |
|---------------------|---------|---|
| capacity_shared | boolean | Specifies whether the capacities are shared across all objects that use this QoS policy-group. Default is false. |
| max_throughput_iops | integer | Maximum throughput defined by this policy. It is specified in terms of IOPS. 0 means no maximum throughput is enforced. |
| max_throughput_mbps | integer | Maximum throughput defined by this policy. It is specified in terms of Mbps. 0 means no maximum throughput is enforced. |
| min_throughput_iops | integer | Minimum throughput defined by this policy. It is specified in terms of IOPS. 0 means no minimum throughput is enforced. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set. |
| min_throughput_mbps | integer | Minimum throughput defined by this policy. It is specified in terms of Mbps. 0 means no minimum throughput is enforced. |

svm

SVM, applies only to SVM-scoped objects.

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

| Name | Type | Description |
|------|--------|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

qos_policy

| Name | Type | Description |
|--------------|--------------------------|---|
| _links | _links | |
| adaptive | adaptive | Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space. |
| fixed | fixed | QoS policy-groups define a fixed service level objective (SLO) for a storage object. |
| name | string | Name of the QoS policy. |
| object_count | integer | Number of objects attached to this policy. |
| pgid | integer | Policy group ID of the QoS policy. |
| policy_class | string | Class of the QoS policy. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | |

job_link

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

error_arguments

| Name | Type | Description |
|---------|--------|------------------|
| code | string | Argument code |
| message | string | Message argument |

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. |

Delete a QoS policy

DELETE /storage/qos/policies/{uuid}

Introduced In: 9.6

Deletes a QoS policy. All QoS workloads associated with the policy are removed.

Related ONTAP commands

- qos policy-group delete
- qos adaptive-policy-group delete

Parameters

| Name | Type | In | Required | Description |
|----------------|---------|-------|----------|--|
| uuid | string | path | True | <ul style="list-style-type: none"> Introduced in: 9.8 |
| return_timeout | integer | query | False | <p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> Default value: 1 Max value: 120 Min value: 0 |

Response

Status: 200, Ok

Error

Status: Default, Error

| 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

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 QoS policy

GET /storage/qos/policies/{uuid}

Introduced In: 9.6

Retrieves a specific QoS policy.

Related ONTAP commands

- `qos policy-group show`
- `qos adaptive-policy-group show`

Parameters

| Name | Type | In | Required | Description |
|--------|---------------|-------|----------|--|
| uuid | string | path | True | <ul style="list-style-type: none">• Introduced in: 9.8 |
| fields | array[string] | query | False | Specify the fields to return. |

Response

Status: 200, Ok

| Name | Type | Description |
|------------------------|--------------------------|---|
| _links | _links | |
| adaptive | adaptive | Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space. |
| fixed | fixed | QoS policy-groups define a fixed service level objective (SLO) for a storage object. |
| name | string | Name of the QoS policy. |
| object_count | integer | Number of objects attached to this policy. |
| pgid | integer | Policy group ID of the QoS policy. |

| Name | Type | Description |
|--------------|--------|---|
| policy_class | string | Class of the QoS policy. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | |

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "adaptive": {
    "block_size": "any",
    "expected_iops_allocation": "used_space",
    "peak_iops_allocation": "used_space"
  },
  "name": "extreme",
  "object_count": 0,
  "pgid": 0,
  "policy_class": "undefined",
  "scope": "cluster",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Error

Status: Default, Error

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

adaptive

Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

| Name | Type | Description |
|--------------------------|---------|---|
| absolute_min_iops | integer | Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set. |
| block_size | string | Specifies the block size |
| expected_iops | integer | Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set. |
| expected_iops_allocation | string | Specifies the size to be used to calculate expected IOPS per TB. The size options are either the storage object allocated space or the storage object used space. |
| peak_iops | integer | Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size. |

| Name | Type | Description |
|----------------------|--------|---|
| peak_iops_allocation | string | Specifies the size to be used to calculate peak IOPS per TB. The size options are either the storage object allocated space or the storage object used space. |

fixed

QoS policy-groups define a fixed service level objective (SLO) for a storage object.

| Name | Type | Description |
|---------------------|---------|---|
| capacity_shared | boolean | Specifies whether the capacities are shared across all objects that use this QoS policy-group. Default is false. |
| max_throughput_iops | integer | Maximum throughput defined by this policy. It is specified in terms of IOPS. 0 means no maximum throughput is enforced. |
| max_throughput_mbps | integer | Maximum throughput defined by this policy. It is specified in terms of Mbps. 0 means no maximum throughput is enforced. |
| min_throughput_iops | integer | Minimum throughput defined by this policy. It is specified in terms of IOPS. 0 means no minimum throughput is enforced. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set. |
| min_throughput_mbps | integer | Minimum throughput defined by this policy. It is specified in terms of Mbps. 0 means no minimum throughput is enforced. |

svm

SVM, applies only to SVM-scoped objects.

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

| Name | Type | Description |
|------|--------|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

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. |

Update a QoS policy

PATCH /storage/qos/policies/{uuid}

Introduced In: 9.6

Update a specific QoS policy.

Related ONTAP commands

- qos policy-group modify
- qos adaptive-policy-group modify

Parameters

| Name | Type | In | Required | Description |
|----------------|---------|-------|----------|--|
| uuid | string | path | True | <ul style="list-style-type: none"> Introduced in: 9.8 |
| return_timeout | integer | query | False | <p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> Default value: 1 Max value: 120 Min value: 0 |

Request Body

| Name | Type | Description |
|----------|--------------------------|---|
| _links | _links | |
| adaptive | adaptive | Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space. |

| Name | Type | Description |
|--------------|---------|---|
| fixed | fixed | QoS policy-groups define a fixed service level objective (SLO) for a storage object. |
| name | string | Name of the QoS policy. |
| object_count | integer | Number of objects attached to this policy. |
| pgid | integer | Policy group ID of the QoS policy. |
| policy_class | string | Class of the QoS policy. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | |

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "adaptive": {
    "block_size": "any",
    "expected_iops_allocation": "used_space",
    "peak_iops_allocation": "used_space"
  },
  "name": "extreme",
  "object_count": 0,
  "pgid": 0,
  "policy_class": "undefined",
  "scope": "cluster",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 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 |
|------------|--|
| 8454147 | The maximum limit for QoS policies has been reached. |
| 8454154 | The name specified for creating conflicts with an existing QoS policy name. |
| 8454260 | Invalid value for maximum and minimum fields. Valid values for max_throughput_iops and max_throughput_mbps combination is for the ratio of max_throughput_mbps and max_throughput_iops to be within 1 to 4096. |
| 8454273 | Invalid value for an adaptive field. Value should be non-zero. |
| 8454277 | The name specified for creating an adaptive QoS policy conflicts with an existing fixed QoS policy name. |
| 8454278 | The name specified for creating a fixed QoS policy conflicts with an existing adaptive QoS policy name. |

| Error Code | Description |
|------------|---|
| 8454286 | Modifications on these cluster scoped preset policies is prohibited. |
| 8454327 | The existing fixed QoS policy cannot be modified to an adaptive QoS policy. |
| 8454328 | The existing adaptive QoS policy cannot be modified to a fixed QoS policy. |

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 | |

adaptive

Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

| Name | Type | Description |
|--------------------------|---------|---|
| absolute_min_iops | integer | Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set. |
| block_size | string | Specifies the block size |
| expected_iops | integer | Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set. |
| expected_iops_allocation | string | Specifies the size to be used to calculate expected IOPS per TB. The size options are either the storage object allocated space or the storage object used space. |
| peak_iops | integer | Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size. |

| Name | Type | Description |
|----------------------|--------|---|
| peak_iops_allocation | string | Specifies the size to be used to calculate peak IOPS per TB. The size options are either the storage object allocated space or the storage object used space. |

fixed

QoS policy-groups define a fixed service level objective (SLO) for a storage object.

| Name | Type | Description |
|---------------------|---------|---|
| capacity_shared | boolean | Specifies whether the capacities are shared across all objects that use this QoS policy-group. Default is false. |
| max_throughput_iops | integer | Maximum throughput defined by this policy. It is specified in terms of IOPS. 0 means no maximum throughput is enforced. |
| max_throughput_mbps | integer | Maximum throughput defined by this policy. It is specified in terms of Mbps. 0 means no maximum throughput is enforced. |
| min_throughput_iops | integer | Minimum throughput defined by this policy. It is specified in terms of IOPS. 0 means no minimum throughput is enforced. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set. |
| min_throughput_mbps | integer | Minimum throughput defined by this policy. It is specified in terms of Mbps. 0 means no minimum throughput is enforced. |

svm

SVM, applies only to SVM-scoped objects.

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

| Name | Type | Description |
|------|--------|---|
| name | string | The name of the SVM. This field cannot be specified in a PATCH method. |
| uuid | string | The unique identifier of the SVM. This field cannot be specified in a PATCH method. |

qos_policy

| Name | Type | Description |
|--------------|--------------------------|---|
| _links | _links | |
| adaptive | adaptive | Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space. |
| fixed | fixed | QoS policy-groups define a fixed service level objective (SLO) for a storage object. |
| name | string | Name of the QoS policy. |
| object_count | integer | Number of objects attached to this policy. |
| pgid | integer | Policy group ID of the QoS policy. |
| policy_class | string | Class of the QoS policy. |
| scope | string | Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects. |
| svm | svm | SVM, applies only to SVM-scoped objects. |
| uuid | string | |

job_link

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

error_arguments

| Name | Type | Description |
|---------|--------|------------------|
| code | string | Argument code |
| message | string | Message argument |

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. |

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.