



qos commands

ONTAP 9.9.1 commands

NetApp

February 12, 2024

Table of Contents

qos commands	1
qos adaptive-policy-group commands	1
qos policy-group commands	7
qos settings commands	13
qos statistics commands	16
qos workload commands	56

qos commands

qos adaptive-policy-group commands

qos adaptive-policy-group create

Create an adaptive policy group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos adaptive-policy-group create` command creates a new adaptive policy group. After the adaptive policy group is created, you can assign one or more storage objects to the policy. When a storage object is assigned to an adaptive policy group, the maximum throughput QoS setting automatically adjusts based on the storage object used space or the storage object allocated space. QoS minimum throughput setting is calculated from the `expected-iops` parameter and the storage object allocated size. It is set only for the storage objects that reside on AFF platforms.

After you create an adaptive policy group, use the `storage object create` command or `storage object modify` command to apply the adaptive policy group to a storage object.

Parameters

-policy-group <text> - Name

Specifies the name of the adaptive policy group. Adaptive policy group names must be unique and are restricted to 127 alphanumeric characters including underscores "_" and hyphens "-". Adaptive policy group names must start with an alphanumeric character. Use the `qos adaptive-policy-group rename` command to change the adaptive policy group name.

-vserver <vserver name> - Vserver

Specifies the data Vserver to which this adaptive policy group belongs to. You can apply this adaptive policy group to only the storage objects contained in the specified Vserver. If the system has only one Vserver, then the command uses that Vserver by default.

-expected-iops {<integer>[IOPS/{GB|TB}]} (default: TB) - Expected IOPS

Specifies the minimum expected IOPS per TB or GB allocated based on the storage object allocated size.

-peak-iops {<integer>[IOPS/{GB|TB}]} (default: TB) - Peak IOPS

Specifies the maximum possible IOPS per TB or GB allocated based on the storage object allocated size or the storage object used size.

[-absolute-min-iops <qos_tput>] - Absolute Minimum IOPS

Specifies the absolute minimum IOPS which is used as an override when the expected IOPS is less than this value. The default value is computed as follows:

if `expected-iops` \geq 6144/TB, then `absolute-min-iops` = 1000IOPS; if `expected-iops` \geq 2048/TB and `expected-iops` $<$ 6144/TB, then `absolute-min-iops` = 500IOPS; if `expected-iops` \geq 1/MB and `expected-iops` $<$ 2048/TB, then `absolute-min-iops` = 75IOPS.

[-expected-iops-allocation {used-space|allocated-space}] - Expected IOPS Allocation

Specifies the expected IOPS allocation policy. The allocation policy is either *allocated-space* or *used-space*. When the expected-iops-allocation policy is set to *allocated-space*, the expected IOPS is calculated based on the size of the storage object. When the expected-iops-allocation policy is set to *used-space*, the expected IOPS is calculated based on the amount of data stored in the storage object taking into account storage efficiencies. The default value is *allocated-space*.

[-peak-iops-allocation {used-space|allocated-space}] - Peak IOPS Allocation

Specifies the peak IOPS allocation policy. The allocation policy is either *allocated-space* or *used-space*. When the peak-iops-allocation policy is set to *allocated-space*, the peak IOPS is calculated based on the size of the storage object. When the peak-iops-allocation policy is set to *used-space*, the peak IOPS is calculated based on the amount of data stored in the storage object taking into account storage efficiencies. The default value is *used-space*.

[-block-size {ANY|4K|8K|16K|32K|64K|128K}] - Block Size

Specifies the I/O block size for the QoS adaptive policy group. The default value is "ANY". When block-size of "ANY" is specified, then control is by IOPS. When block-size other than "ANY" is specified, then control is by IOPS and bytes per second(bps). bps is the product of IOPS and block-size.

Examples

```
cluster1::> qos adaptive-policy-group create p1 -vserver vs1  
          -expected-iops 100IOPS/TB -peak-iops 1000/TB
```

Creates the "p1" adaptive policy group which belongs to Vserver "vs1" with expected-iops of 100IOPS/TB and peak-iops of 1000IOPS/TB with default value for absolute-min-iops

```
cluster1::> qos adaptive-policy-group create p2 -vserver vs1  
          -expected-iops 100IOPS/GB -peak-iops 1000IOPS/GB  
          -absolute-min-iops 200IOPS
```

Creates the "p1" adaptive policy group which belongs to Vserver "vs1" with expected-iops of 100IOPS/TB and peak-iops of 1000IOPS/TB with the absolute-min-iops set to 200IOPS.

Related Links

- [qos adaptive-policy-group rename](#)

qos adaptive-policy-group delete

Delete an adaptive policy group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos adaptive-policy-group delete` command deletes an adaptive policy group from a cluster. You cannot delete a policy group if a QoS workload associated with a storage object is assigned to it unless you

use the `-force` parameter. Using the `-force` parameter will delete all the QoS workloads for storage objects associated with the specified adaptive policy groups .

Only user created adaptive policy groups can be deleted. Default adaptive policy groups are read only and cannot be deleted.

Parameters

`-policy-group <text>` - Name

Specifies the name of the adaptive policy group that you want to delete.

`[-force <true>]` - Force Delete Workloads for the QoS adaptive policy group (privilege: advanced)

Specifies whether to delete an adaptive policy group along with any underlying workloads.

Examples

The following example deletes "p1" adaptive policy group:

```
cluster1::> qos adaptive-policy-group delete p1
```

Deletes the "p1" adaptive policy group along with any underlying QoS workloads.

```
cluster1::> qos adaptive-policy-group delete p1 -force
```

qos adaptive-policy-group modify

Modify an adaptive policy group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos adaptive-policy-group modify` command modifies an adaptive policy group.

Only user-created adaptive policy groups can be modified. Default adaptive policy groups are read-only and cannot be modified.

Parameters

`-policy-group <text>` - Name

Specifies the name of the adaptive policy group. Adaptive policy group names must be unique and are restricted to 127 alphanumeric characters including underscores "_" and hyphens "-". Adaptive policy group names must start with an alphanumeric character. Use the [qos adaptive-policy-group rename](#) command to change the adaptive policy group name.

`[-expected-iops {<integer>}[IOPS / {GB | TB}]] (default: TB)}` - Expected IOPS

Specifies the minimum expected IOPS per TB or GB allocated based on the storage object allocated size.

QoS minimum throughput setting is calculated from the expected-iops parameter. It is set only for the storage objects that reside on AFF platforms.

[-peak-iops {<integer>}[IOPS / {GB|TB}]] (default: TB)] - Peak IOPS

Specifies the maximum possible IOPS per TB or GB allocated based on the storage object allocated size or the storage object used size.

[-absolute-min-iops <qos_tput>] - Absolute Minimum IOPS

Specifies the absolute minimum IOPS which is used as an override when the expected IOPS is less than this value. The default value is computed as follows:

if expected-iops >= 6144/TB, then absolute-min-iops = 1000IOPS; if expected-iops >= 2048/TB and expected-iops < 6144/TB, then absolute-min-iops = 500IOPS; if expected-iops >= 1/MB and expected-iops < 2048/TB, then absolute-min-iops = 75IOPS.

[-expected-iops-allocation {used-space|allocated-space}] - Expected IOPS Allocation

Specifies the expected IOPS allocation policy. The allocation policy is either *allocated-space* or *used-space*. When the expected-iops-allocation policy is set to *allocated-space*, the expected IOPS is calculated based on the size of the storage object. When the expected-iops-allocation policy is set to *used-space*, the expected IOPS is calculated based on the amount of data stored in the storage object taking into account storage efficiencies. The default value is *allocated-space*.

[-peak-iops-allocation {used-space|allocated-space}] - Peak IOPS Allocation

Specifies the peak IOPS allocation policy. The allocation policy is either *allocated-space* or *used-space*. When the peak-iops-allocation policy is set to *allocated-space*, the peak IOPS is calculated based on the size of the storage object. When the peak-iops-allocation policy is set to *used-space*, the peak IOPS is calculated based on the amount of data stored in the storage object taking into account storage efficiencies. The default value is *used-space*.

[-block-size {ANY|4K|8K|16K|32K|64K|128K}] - Block Size

Specifies the I/O block size for the QoS adaptive policy group. The default value is "ANY". When block-size of "ANY" is specified, then control is by IOPS. When block-size other than "ANY" is specified, then control is by IOPS and bytes per second(bps). bps is the product of IOPS and block-size.

Examples

The following example modifies the "p1" adaptive policy group with specified values.

```
cluster1::> qos adaptive-policy-group modify -policy-group p1  
          -expected-iops 200IOPS/TB -peak-iops 2000IOPS/TB  
          -absolute-min-iops 100IOPS
```

Related Links

- [qos adaptive-policy-group rename](#)

qos adaptive-policy-group rename

Rename an adaptive policy group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos adaptive-policy-group rename` command changes the name of an existing adaptive policy group.

Parameters

-policy-group <text> - Name

Specifies the existing name of the adaptive policy group that you want to rename.

-new-name <text> - New adaptive policy group name

Specifies the new name of the adaptive policy group. Adaptive policy group names must be unique and are restricted to 127 alphanumeric characters including underscores "_" and hyphens "-". Adaptive policy group names must start with an alphanumeric character.

Examples

```
cluster1::> qos adaptive-policy-group rename -policy-group p1 -new-name p1_new
```

Renames the adaptive policy group from "p1" to "p1_new".

qos adaptive-policy-group show

Display a list of adaptive policy groups

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos adaptive-policy-group show` command shows the current settings of the adaptive policy groups on a cluster. You can view the list of adaptive policy groups and also the detailed information about a specific adaptive policy group.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use '`-fields ?`' to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-policy-group <text>] - Name

Selects the adaptive policy groups that match this parameter value.

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

[-vserver <vserver name>] - Vserver

If this parameter is specified, the command displays information only about the adaptive policy groups with a matching vserver.

[-uuid <UUID>] - Uuid

If this parameter is specified, the command displays information only about the adaptive policy groups with the specified UUID.

[-pgid <integer>] - ID

If this parameter is specified, the command displays information only about the adaptive policy groups that match the given policy group ID, which is an integer that uniquely identifies the adaptive policy group.

[-expected-iops {<integer>[IOPS / {GB|TB}]} (default: TB)] - Expected IOPS

If this parameter is specified, the command displays information only about the adaptive policy groups with the specified minimum expected IOPS per TB or GB.

[-peak-iops {<integer>[IOPS / {GB|TB}]} (default: TB)] - Peak IOPS

If this parameter is specified, the command displays information only about the adaptive policy groups with the specified maximum possible IOPS per TB or GB.

[-absolute-min-iops <qos_tput>] - Absolute Minimum IOPS

If this parameter is specified, the command displays information only about the adaptive policy groups with the specified absolute minimum IOPS.

[-expected-iops-allocation {used-space|allocated-space}] - Expected IOPS Allocation

If this parameter is specified, the command displays information only about the adaptive policy groups with the specified expected IOPS allocation policy used to compute the expected IOPS per TB or GB.

[-peak-iops-allocation {used-space|allocated-space}] - Peak IOPS Allocation

If this parameter is specified, the command displays information only about the adaptive policy groups with the specified peak IOPS allocation policy used to compute the maximum possible IOPS per TB or GB.

[-block-size {ANY|4K|8K|16K|32K|64K|128K}] - Block Size

Specifies the I/O block size for the QoS adaptive policy group. The default value is "ANY". When block-size of "ANY" is specified, then control is by IOPS. When block-size other than "ANY" is specified, then control is by IOPS and bytes per second(bps). bps is the product of IOPS and block-size.

[-num-workloads <integer>] - Number of Workloads

If this parameter is specified, the command displays information only about the adaptive policy groups with the specified number of workloads.

Examples

The example above displays all adaptive policy groups on the cluster.

```

cluster1::> qos adaptive-policy-group show
qos adaptive-policy-group show
Expected      Peak          Minimum     Block
Name          Vserver   Wklds    IOPS        IOPS        IOPS        Size
-----
apg1           vs1       1       100IOPS/MB  1000IOPS/MB  75IOPS     8K
apg2           vs1       1       100IOPS/MB  1000IOPS/MB  75IOPS     4K
extreme        clus-1   0       6144IOPS/TB 12288IOPS/TB 1000IOPS  ANY
performance    clus-1   0       2048IOPS/TB  4096IOPS/TB  500IOPS  ANY
value          clus-1   0       128IOPS/TB   512IOPS/TB   75IOPS  ANY
5 entries were displayed.

```

qos policy-group commands

qos policy-group create

Create a policy group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos policy-group create` command creates a new policy group. You can use a QoS policy group to control a set of storage objects known as "workloads" - LUNs, volumes, files, or Vservers. Policy groups define measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated.

After you create a policy group, you use the storage object `create` command or the storage object `modify` command to apply the policy group to a storage object.

Parameters

-policy-group <text> - Policy Group Name

Specifies the name of the policy group. Policy group names must be unique and are restricted to 127 alphanumeric characters including underscores "_" and hyphens "-". Policy group names must start with an alphanumeric character. You use the [qos policy-group rename](#) command to change the policy group name.

-vserver <vserver name> - Vserver

Specifies the data Vserver to which this policy group belongs. You can apply this policy group to only the storage objects contained in the specified Vserver. For example, if you want to apply this policy group to a volume, that volume must belong to the specified Vserver. Using this parameter does not apply the policy group's SLOs to the Vserver. You need to use the `vserver modify` command if you want to apply this policy group to the Vserver. If the system has only one Vserver, then the command uses that Vserver by default.

[-max-throughput <qos_tput>] - Maximum Throughput

Specifies the maximum throughput for the policy group. A maximum throughput limit specifies the throughput that the policy group must not exceed. It is specified in terms of IOPS or MB/s, or a combination of comma separated IOPS and MB/s. The range is one to infinity. A value of zero is accepted but is

internally treated as infinity.

The values entered here are case-insensitive, and the units are base ten. There should be no space between the number and the units. The default value for max-throughput is infinity, which can be specified by the special value "INF". Note that there is no default unit - all numbers except zero require explicit specification of the units.

Two reserved keywords, "none" and "INF", are available for the situation that requires removal of a value, and the situation that needs to specify the maximum available value.

Examples of valid throughput specifications are: "100B/s", "10KB/s", "1gb/s", "500MB/s", "1tb/s", "100iops", "100iops,400KB/s", and "800KB/s,100iops"

[-min-throughput <qos_tput>] - Minimum Throughput

Specifies the minimum throughput for the policy group. A minimum throughput specifies the desired performance level for a policy group. It is specified in terms of IOPS or MB/s, or a combination of comma separated IOPS and MB/s.

The values entered here are case-insensitive, and the units are base ten. There should be no space between the number and the units. The default value for min-throughput is "0". The default unit is IOPS.

One reserved keyword, 'none' is available for the situation that requires removal of a value.

Examples of valid throughput specifications are: "100B/s", "10KB/s", "1gb/s", "500MB/s", "1tb/s", "100iops", "100iops,400KB/s", and "800KB/s,100iops"

[-is-shared {true|false}] - Is Shared

Specifies whether the policy group can be shared or not. The default value is "true". This parameter specifies if the SLOs of the policy group are shared between the workloads or if the SLOs are applied separately to each workload.

Examples

```
cluster1::> qos policy-group create p1 -vserver vs1
```

Creates the "p1" policy group which belongs to Vserver "vs1" with default policy values.

```
cluster1::> qos policy-group create p2 -vserver vs1 -max-throughput  
500MB/s
```

Creates the "p2" policy group which belongs to Vserver "vs1" with the maximum throughput set to 500 MB/s.

```
cluster1::> qos policy-group create p3 -vserver vs1 -max-throughput  
500MB/s -is-shared false
```

Creates the "p3" policy group which belongs to Vserver "vs1" with the maximum throughput set to 500 MB/s and shared set to false.

Related Links

- [qos policy-group rename](#)

qos policy-group delete

Delete an existing QoS Policy Group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos policy-group delete` command deletes a policy group from a cluster. You cannot delete a policy group if a qos workload associated with storage object is assigned to it unless you use "-force". Using "-force" will delete all the qos workloads for storage objects associated with the specified policy groups .

You can only delete user-defined policy groups. You cannot delete preset policy groups.

Parameters

-policy-group <text> - Policy Group Name

Specifies the name of the policy group that you want to delete.

[-force <true>] - Force Delete Workloads for the QoS Policy Group (privilege: advanced)

Specifies whether to delete a policy group along with any underlying workloads.

Examples

```
cluster1::> qos policy-group delete p1
```

Deletes the "p1" policy group.

```
cluster1::> qos policy-group delete p1 -force
```

Deletes the "p1" policy group along with any underlying qos workloads.

qos policy-group modify

Modify a policy group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos policy-group modify` command modifies a user-created policy group.

Parameters

-policy-group <text> - Policy Group Name

Specifies the name of the policy group that you want to modify.

[-max-throughput <qos_tput>] - Maximum Throughput

Specifies the maximum throughput for the policy group. A maximum throughput limit specifies the throughput that the policy group must not exceed. It is specified in terms of IOPS or MB/s, or a combination of comma separated IOPS and MB/s. The range is one to infinity. A value of zero is accepted but is internally treated as infinity.

The values entered here are case-insensitive, and the units are base ten. There should be no space between the number and the units. The default value for max-throughput is infinity, which can be specified by the special value "INF". Note there is no default unit - all numbers except zero require explicit specification of the units.

Two reserved keywords, "none" and "INF", are available for the situation that requires removal of a value, and the situation that needs to specify the maximum available value.

Examples of valid throughput specifications are: "100B/s", "10KB/s", "1gb/s", "500MB/s", "1tb/s", and "100iops".

[-min-throughput <qos_tput>] - Minimum Throughput

Specifies the minimum throughput for the policy group. A minimum throughput specifies the desired performance level for a policy group. It is specified in terms of IOPS or MB/s, or a combination of comma separated IOPS and MB/s.

The values entered here are case-insensitive, and the units are base ten. There should be no space between the number and the units. The default value for min-throughput is "0". The default unit is IOPS.

One reserved keyword, 'none' is available for the situation that requires removal of a value.

Examples of valid throughput specifications are: "100B/s", "10KB/s", "1gb/s", "500MB/s", "1tb/s", "100iops", "100iops,400KB/s", and "800KB/s,100iops"

Examples

```
cluster1::> qos policy-group modify p1 -max-throughput 10IOPS
```

Modifies the "p1" policy group and sets its max throughput value to 10 IOPS.

qos policy-group rename

Rename a policy group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos policy-group rename` command changes the name of an existing policy group.

Parameters

-policy-group <text> - Policy Group Name

Specifies the existing name of the policy group that you want to rename.

-new-name <text> - New Policy Group Name

Specifies the new name of the policy group. Policy group names must be unique and are restricted to 127 alphanumeric characters including underscores "_" and hyphens "-". Policy group names must start with an alphanumeric character.

Examples

```
cluster1::> qos policy-group rename -policy-group p1 -new-name p1_new
```

Renames the policy group from "p1" to "p1_new".

qos policy-group show

Display a list of policy groups

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos policy-group show` command shows the current settings of the policy groups on a cluster. You can display a list of the policy groups and you can view detailed information about a specific policy group.

Parameters

{ [-fields <fieldname>, ...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use '`-fields ?`' to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-policy-group <text>] - Policy Group Name

Selects the policy groups that match this parameter value

Policy groups define measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated.

[-vserver <vserver name>] - Vserver

Selects the policy groups that match this parameter value

[-uuid <UUID>] - Uuid

Selects the policy groups that match this parameter value

`[-class <QoS Configuration Class>] - Policy Group Class`

Selects the policy groups that match this parameter value

`[-pgid <integer>] - Policy Group ID`

Selects the policy groups that match this parameter value

This uniquely identifies the policy group

`[-max-throughput <qos_tput>] - Maximum Throughput`

Selects the policy groups that match this parameter value

A maximum throughput limit specifies the throughput (in IOPS or MB/s) that the policy group must not exceed.

`[-min-throughput <qos_tput>] - Minimum Throughput`

Selects the policy groups that match this parameter value

A minimum throughput specifies the desired performance level for a policy group.

`[-num-workloads <integer>] - Number of Workloads`

Selects the policy groups that match this parameter value.

`[-throughput-policy <text>] - Throughput Policy`

Selects the policy groups that match this parameter value. You can specify the throughput range in terms of IOPS or data rate. For example, 0-INF, 0-400IOPS, 0-200KB/s, 0-400MB/s .

`[-is-shared {true|false}] - Is Shared`

Selects the policy groups that match this parameter value.

The shared value specifies whether the policy group is a shared policy group or not.

`[-is-auto-generated {true|false}] - Is Policy Auto Generated`

Selects the policy groups that match this parameter value.

The auto-generated value specifies whether the policy group is an automatically generated policy group or not.

Examples

```
cluster1::> qos policy-group show
Name          Vserver      Class        Wklds  Throughput
-----  -----
pg1           vs4         user-defined 0      0-200IOPS
pg2           vs0         user-defined 0      0-500IOPS
pg5           vs0         user-defined 0      0-300IOPS
pg6           vs0         user-defined 0      0-INF
4 entries were displayed.
```

The example above displays all policy groups on the cluster.

qos settings commands

qos settings throughput-floors-v2

Enable/Disable throughput floors v2 on AFF

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `qos settings throughput-floors-v2` command is used to enable or disable floors v2 on AFF. The default is enabled. With floors v2 enabled, throughput floors can be met when controllers are heavily utilized at the expense of higher latency on other workloads. Floors v2 applies to both QoS and Adaptive QoS.

Parameters

-enable {true|false} - enable or disable throughput floors v2 on AFF (privilege: advanced)

This specifies if floors v2 is enabled or disabled. If this parameter is specified with `false` floors v2 will be disabled.

Examples

The following example disables floors v2 on AFF.

```
cluster1::*> qos settings throughput-floors-v2 -enable false
```

qos settings cache modify

Modify the cache policy

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos settings cache modify` command modifies the existing default caching-policy. The list of caching policies can be obtained from the `qos setting cache show -fields cache-setting` command.

Parameters

-cache-setting <text> - Cache Policy Name

Valid inputs to this parameter include any one of the listed caching-policies. This command is to be used together with the `-default` parameter. If you use this parameter, the command modifies the specified caching-policy based on the `-default` parameter.

[-default {true|false}] - Is Default?

Valid inputs to this parameter are true and false. Together with `cache-setting`, this parameter helps set or unset a caching-policy as default.

Examples

```
cluster1::> qos settings cache modify -default true -cache-setting  
random_read_write-random_write
```

Sets caching-policy random_read_write-random_write as default.

qos settings cache show

Display list of cache policies

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The ` qos settings cache show` shows the current caching-policies, class to which they belong, the number of workloads associated with each of the policies, and whether or not they are set to default. The following external-cache policies are available:

- none - Does not cache any user data or metadata blocks.
- auto - Read caches all metadata and randomly read user data blocks, and write caches all randomly overwritten user data blocks.
- all - Read caches all data blocks read and written. It does not do any write caching.
- all-random_write - Read caches all data blocks read and written. It also write caches randomly overwritten user data blocks.
- all_read - Read caches all metadata, randomly read, and sequentially read user data blocks.
- all_read-random_write - Read caches all metadata, randomly read, and sequentially read user data blocks. It also write caches randomly overwritten user data blocks.
- all_read_random_write - Read caches all metadata, randomly read, sequentially read and randomly written user data.
- all_read_random_write-random_write - Read caches all metadata, randomly read, sequentially read, and randomly written user data blocks. It also write caches randomly overwritten user data blocks.
- meta - Read caches only metadata blocks.
- meta-random_write - Read caches all metadata and write caches randomly overwritten user data blocks.
- noread-random_write - Write caches all randomly overwritten user data blocks. It does not do any read caching.
- random_read - Read caches all metadata and randomly read user data blocks.
- random_read_write - Read caches all metadata, randomly read and randomly written user data blocks.
- random_read_write-random_write - Read caches all metadata, randomly read, and randomly written user data blocks. It also write caches randomly overwritten user data blocks.



Note that in a caching-policy name, a hyphen (-) separates read and write caching policies.

Parameters

{ [-fields <fieldname>, ...]

The input to this parameter is one of the following: {cache-setting|class|default|num-workloads}. If you use this parameter, the command displays information related to the specified input field.

| [-instance] }

If you use this parameter, the command displays information about the caching-policies in a list format.

[-cache-setting <text>] - Cache Policy Name

The input to this parameter is any one of the above listed caching-policies. If you use this parameter, the command displays information corresponding to the specified caching-policy.

[-class <QoS Configuration Class>] - Cache Policy Class

The input to this parameter is one of the following: {undefined|preset|user-defined|system-defined|autovolume}. If you use this parameter, the command displays information corresponding to the specified policy-group class.

[-default {true|false}] - Is Default?

The input to this parameter is true and false. If you use this parameter, the command displays information corresponding to entries that have the specified default value.

[-num-workloads <integer>] - Number Of Workloads With This Policy

The input to this parameter is an integer. If you use this parameter, the command displays information about policy-groups matching the specified number of workloads.

Examples

```

cluster1::> qos settings cache show
Policy Name    Class      Num Workloads   Default
-----  -----  -----  -----
all          preset      0            false
all-random_write
            preset      0            false
all_read      preset      0            false
all_read-random_write
            preset      0            false
all_read_random_write
            preset      0            false
all_read_random_write-random_write
            preset      0            false
auto         preset      2            false
meta         preset      0            false
meta-random_write
            preset      0            false
none         preset      0            false
noread-random_write
            preset      0            false
random_read  preset     25           false
random_read_write
            preset      0            false
random_read_write-random_write
            preset     28           true
14 entries were displayed.

```

Shows QoS settings for the caching policies.

qos statistics commands

qos statistics characteristics show

Display QoS policy group characterization

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos statistics characteristics show` command displays data that characterizes the behavior of QoS policy groups.

The command displays the following data:

- The QoS policy group name (Policy Group)
- Input/output operations performed per second (IOPS)

- Throughput achieved in kilobytes per second (KB/s) or megabytes per second (MB/s) as appropriate (Throughput)
- Request size in bytes (B) (Request size)
- Read percentage from total I/O (Read)
- Concurrency, which indicates the number of concurrent users generating the I/O traffic (Concurrency)

The results displayed per iteration are sorted by IOPS. Each iteration starts with a row that displays the total IOPS used across all QoS policy groups. Other columns in this row are either totals or averages.

Parameters

[-node {<nodename>|local}] - Node

Selects the policy groups that match this parameter value. If you do not specify this parameter, the command displays data for the entire cluster.

[-iterations <integer>] - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

{ [-rows <integer>] - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. Valid values are from 1 to 20. The default value is 10.

| [-policy-group <text>] - QoS Policy Group Name

Selects the QoS policy group whose name matches the specified value. If you do not specify this parameter, the command displays data for all QoS policy groups.

| [-adaptive-policy-group <text>] - Adaptive QoS Policy Group Name }

Selects the QoS adaptive policy group whose name matches the specified value. If you do not specify this parameter, the command displays data for all QoS policy groups.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

Examples

```

cluster1::> qos statistics characteristics show -iterations 100 -rows 4
Policy Group          IOPS      Throughput Request size   Read
Concurrency
-----
-----  

-total-                31       304.00KB/s     10041B    0%
16
_System-Best-Effort  15       0KB/s        0B        0%
0
vol1                  11       44.00KB/s     4096B    0%
40
vol2                  4        256.00KB/s    65536B    0%
14
vs1vol0               1        4.00KB/s     4096B    0%
4
-total-                37       808.00KB/s    22361B    2%
3
_System-Best-Effort  15       0KB/s        0B        0%
0
vol2                  12       768.00KB/s    65536B    0%
9
vs1vol0               8        32.00KB/s     4096B    12%
1
vol1                  2        8.00KB/s      4096B    0%
1

```

The example above displays the characteristics of the 4 QoS policy groups with the highest IOPS values and refreshes the display 100 times before terminating.

```

cluster1::> qos statistics characteristics show -iterations 100 -policy
-group pg1
Policy Group          IOPS      Throughput Request size Read
Concurrency
-----
-----  

-total-                293       3.02MB/s    10783B  54%
0
pg1                   118       470.67KB/s   4096B  100%
0
-total-                181       478.14KB/s   2700B  65%
0
pg1                   117       469.33KB/s   4096B  100%
0
-total-                226       525.78KB/s   2382B  60%
1
pg1                   110       440.00KB/s   4096B  100%
1
-total-                233       1.67MB/s     7527B  49%
1
pg1                   112       446.67KB/s   4096B  100%
1

```

The example above displays the system characteristics of the QoS policy group *pg1* and refreshes the display 100 times before terminating.

qos statistics latency show

Display latency breakdown data per QoS policy group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos statistics latency show` command displays the average latencies for QoS policy groups across the various Data ONTAP subsystems.

The command displays the following data:

- The QoS policy group name (Policy Group)
- Total latency observed per I/O operation (Latency)
- Latency observed per I/O operation in the Network subsystem (Network)
- Latency observed per I/O operation across the internally connected nodes in a Cluster (Cluster)
- Latency observed per I/O operation in the Data management subsystem (Data)
- Latency observed per I/O operation in the Storage subsystem (Disk)

- Latency observed per I/O operation for QoS Policy Group Ceiling (QoS Max)
 - Latency observed per I/O operation for QoS Policy Group Floor (QoS Min)
 - Latency observed per I/O operation for NVRAM transfer (NVRAM)
 - Latency observed per I/O operation for Object Store (Cloud) operations
 - Latency observed per I/O operation for FlexCache (FlexCache) operations
 - Latency observed per I/O operation for Synchronous Snapmirror (SM Sync) operations
 - Latency observed per I/O operation for Volume Activation (VA) operations

The results displayed per iteration are sorted by the Latency field. Each iteration starts with a row that displays the average latency, in microseconds (us) or milliseconds (ms), observed across all QoS policy groups.

Parameters

[-node {<nodename>|local}] - Node

Selects the policy groups that match this parameter value. If you do not specify this parameter, the command displays data for the entire cluster.

[-iterations <integer>] - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

{ [-rows <integer>] - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. Valid values are from 1 to 20. The default value is 10.

[-policy-group <text>] - QoS Policy Group Name

Selects the QoS policy group whose name matches the specified value. If you do not specify this parameter, the command displays data for all QoS policy groups.

[-adaptive-policy-group <text>] - Adaptive QoS Policy Group Name }

Selects the QoS adaptive policy group whose name matches the specified value. If you do not specify this parameter, the command displays data for all QoS policy groups.

`[-refresh-display {true|false}]` - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

Examples

```
cluster1::> qos statistics latency show -iterations 100 -rows 3
Policy Group          Latency      Network     Cluster      Data
Disk    QoS Max      QoS Min      NVRAM       Cloud      FlexCache   SM Sync
VA
-----
-----
```

0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
vs1vol0		167.82ms	167.22ms		0ms	603.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
vol1		117.76ms	117.56ms		0ms	191.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
vol2		44.24ms	44.05ms		0ms	190.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
-total-		38.89ms	38.63ms		0ms	256.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
vol2		64.47ms	64.20ms		0ms	266.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
vol1		27.28ms	27.03ms		0ms	253.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
vs1vol0		23.72ms	23.47ms		0ms	249.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
-total-		409.81ms	409.65ms		0ms	169.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
vol1		816.92ms	816.80ms		0ms	120.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
vol2		407.88ms	407.66ms		0ms	219.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
vs1vol0		3.68ms	3.49ms		0ms	193.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
-total-		1169.00us	107.00us		0ms	1062.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						
vol2		1169.00us	107.00us		0ms	1062.00us
0ms	0ms	0ms	0ms	0ms	0ms	0ms
0ms						

The example above displays latencies for the 3 QoS policy groups with the highest latencies and refreshes the display 100 times before terminating.

cluster1::> qos statistics latency show -iterations 100 -policy-group pg1							
Policy Group		Latency		Network		Cluster	
Disk	QoS	Max	QoS	Min	NVRAM	Cloud	FlexCache
VA							
-total-			5.88ms	308.00us		0ms	434.00us
5.14ms		0ms	0ms	0ms		0ms	0ms
0ms							
pg1			5.88ms	308.00us		0ms	434.00us
5.14ms		0ms	0ms	0ms		0ms	0ms
0ms							
-total-			4.17ms	280.00us		0ms	477.00us
3.42ms		0ms	0ms	0ms		0ms	0ms
0ms							
pg1			4.17ms	280.00us		0ms	477.00us
3.42ms		0ms	0ms	0ms		0ms	0ms
0ms							
-total-			4.43ms	274.00us		0ms	656.00us
3.50ms		0ms	0ms	0ms		0ms	0ms
0ms							
pg1			4.43ms	274.00us		0ms	656.00us
3.50ms		0ms	0ms	0ms		0ms	0ms
0ms							
-total-			4.89ms	276.00us		0ms	699.00us
3.92ms		0ms	0ms	0ms		0ms	0ms
0ms							
pg1			4.89ms	276.00us		0ms	699.00us
3.92ms		0ms	0ms	0ms		0ms	0ms
0ms							

The example above displays latencies for the QoS policy group *pg1* and refreshes the display 100 times before terminating.

qos statistics performance show

Display system performance data per QoS policy group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos statistics performance show` command shows the current system performance levels that QoS policy groups are achieving.

The command displays the following data:

- The QoS policy group name (Policy Group)
- Input/output operations performed per second (IOPS)
- Throughput in kilobytes per second (KB/s) or megabytes per second (MB/s) as appropriate (Throughput)
- Latency observed per request in microseconds (us) or milliseconds (ms) as appropriate (Latency)

The results displayed per iteration are sorted by IOPS. Each iteration starts with a row that displays the total IOPS used across all QoS policy groups. Other columns in this row are either totals or averages.

Parameters

[-node {<nodename>|local}] - Node

Selects the policy groups that match this parameter value. If you do not specify this parameter, the command displays data for the entire cluster.

[-iterations <integer>] - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

{ [-rows <integer>] - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. Valid values are from 1 to 20. The default value is 10.

| [-policy-group <text>] - QoS Policy Group Name

Selects the QoS policy group whose name matches the specified value. If you do not specify this parameter, the command displays data for all QoS policy groups.

| [-adaptive-policy-group <text>] - Adaptive QoS Policy Group Name }

Selects the QoS adaptive policy group whose name matches the specified value. If you do not specify this parameter, the command displays data for all QoS policy groups.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

Examples

```
cluster1::> qos statistics performance show -iterations 100 -rows 4
Policy Group          IOPS      Throughput     Latency
-----  -----  -----
-total-                79       1296.00KB/s   337.41ms
_System-Best-Effort    25        0KB/s        0ms
vol1                  24       96.00KB/s    193.72ms
vol2                  18      1152.00KB/s   750.98ms
vs1vol0                12       48.00KB/s    707.38ms
-total-                109      1.99MB/s    133.27ms
_System-Best-Effort    35        0KB/s        0ms
vol2                  29       1.81MB/s    249.27ms
vs1vol0                24       96.00KB/s    48.32ms
vol1                  21       84.00KB/s    292.30ms
```

The example above displays the system performance for the 4 QoS policy groups with the highest IOPS and it refreshes the display 100 times before terminating.

```
cluster1::> qos statistics performance show -iterations 100 -policy-group pg1
Policy Group          IOPS      Throughput     Latency
-----  -----  -----
-total-                2833     10.66MB/s   924.00us
pg1                   2655     10.37MB/s   917.00us
-total-                2837     10.65MB/s   923.00us
pg1                   2655     10.37MB/s   917.00us
-total-                2799     10.73MB/s   802.00us
pg1                   2737     10.69MB/s   815.00us
-total-                2930     13.33MB/s   905.00us
pg1                   2720     10.62MB/s   858.00us
```

The example above displays the system performance for the QoS policy group *pg1* and refreshes the display 100 times before terminating.

qos statistics resource cpu show

Display CPU resource utilization data per QoS policy group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos statistics resource cpu show` command displays the CPU utilization for QoS policy groups per node.

The command displays the following data:

- The QoS policy group name (Policy Group)
- CPU utilization observed in percentage (CPU)

The results displayed per iteration are sorted by total CPU utilization. Each iteration starts with a row that displays the total CPU utilization across all QoS policy groups.

Parameters

-node {<nodename>|local} - Node

Selects the policy groups that match this parameter value.

[-iterations <integer>] - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

{ [-rows <integer>] - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. Valid values are from 1 to 20. The default value is 10.

| [-policy-group <text>] - QoS Policy Group Name

Selects the QoS policy group whose name matches the specified value. If you do not specify this parameter, the command displays data for all QoS policy groups.

| [-adaptive-policy-group <text>] - Adaptive QoS Policy Group Name }

Selects the QoS adaptive policy group whose name matches the specified value. If you do not specify this parameter, the command displays data for all QoS policy groups.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

Examples

```
cluster1::> qos statistics resource cpu show -node nodeA -iterations 100
-rows 3
      Policy Group          CPU
      ----- -----
      -total- (100%)        9%
      fast                  1%
      slow                  3%
      medium                5%
      -total- (100%)        8%
      slow                  1%
      fast                  3%
      medium                3%
```

The example above displays the total CPU utilization for the 3 QoS policy groups with the highest CPU utilization and it refreshes the display 100 times before terminating.

```

cluster1::> qos statistics resource cpu show -node local -iterations 100
-priority-group pg1
  Policy Group          CPU
  -----
  -total- (100%)        7%
  pg1                  1%
  -total- (100%)        7%
  pg1                  1%
  -total- (100%)        7%
  pg1                  1%
  -total- (100%)        10%
  pg1                  1%

```

The example above displays the total CPU utilization for the QoS policy group *pg1* and refreshes the display 100 times before terminating.

qos statistics resource disk show

Display disk resource utilization data per QoS policy group

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos statistics resource disk show` command displays the disk utilization for QoS policy groups per node. The disk utilization shows the percentage of time spent on the disk during read and write operations. The command displays disk utilization for system-defined policy groups; however, their disk utilization is not included in the total utilization. The command only supports hard disks.

The command displays the following data:

- The QoS policy group name (Policy Group)
- Disk utilization (Disk)
- The number of HDD data disks utilized (Number of HDD Disks)

The results displayed are sorted by total disk utilization. Each iteration starts with a row that displays the total disk utilization across all QoS policy groups.

Parameters

-node {<nodename>}|local - Node

Selects the policy groups that match this parameter value.

[-iterations <integer>] - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

{ [-rows <integer>] - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. Valid values are from 1 to 20. The default value is 10.

| [-policy-group <text>] - QoS Policy Group Name

Selects the QoS policy group whose name matches the specified value. If you do not specify this parameter, the command displays data for all QoS policy groups.

| [-adaptive-policy-group <text>] - Adaptive QoS Policy Group Name }

Selects the QoS adaptive policy group whose name matches the specified value. If you do not specify this parameter, the command displays data for all QoS policy groups.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

Examples

```
cluster1::> qos statistics resource disk show -node nodeA -iterations 100  
-rows 3  


| Policy Group    | Disk Number | of HDD Disks |
|-----------------|-------------|--------------|
| -total-         | 40%         | 27           |
| pg1             | 22%         | 5            |
| slow            | 10%         | 10           |
| fast            | 8%          | 12           |
| _System_Default | 7%          | 20           |
| -total-         | 42%         | 27           |
| pg1             | 22%         | 5            |
| slow            | 12%         | 10           |
| fast            | 8%          | 12           |
| _System_Default | 7%          | 20           |


```

The example above displays the total disk utilization for the 3 QoS policy groups with the highest disk utilization and it refreshes the display 100 times before terminating.

```

cluster1::> qos statistics resource disk show -node local -iterations 100
-priority-group pg1
      Policy Group          Disk Number of HDD Disks
      -----  -----
      -total-                3%                  10
      pg1                   1%                  24
      -total-                3%                  10
      pg1                   1%                  24
      -total-                3%                  10
      pg1                   1%                  24
      -total-                3%                  10
      pg1                   1%                  24

```

The example above displays the total disk utilization for the QoS policy group *pg1* and refreshes the display 100 times before terminating.

qos statistics volume characteristics show

Display volume characteristics

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos statistics volume characteristics show` command displays data that characterizes the behavior of volumes.

The command displays the following data:

- QoS volume name (Workload)
- QoS workload ID (ID)
- Input/output operations per second (IOPS)
- Throughput achieved in kilobytes per second (KB/s) or megabytes per second (MB/s) as appropriate (Throughput)
- Request size in bytes (B) (Request size)
- Read percentage from total IOPS (Read)
- Concurrency, which indicates the number of concurrent users generating the I/O traffic (Concurrency)

The results displayed per iteration are sorted by IOPS. Each iteration starts with a row that displays the total IOPS used across all volumes. Other columns in this row are either totals or averages.

Parameters

[-node {<nodename>}|local] - Node

Selects the volumes that match this parameter value. If you do not specify this parameter, the command displays data for the entire cluster.

{ [-rows <integer>} - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. The default setting is 10. The allowed range of values is 1 to 20.

| -vserver <vserver name> - Vserver Name

Specifies the Vserver to which the volume belongs.

-volume <volume name> - Volume Name }

Selects the characteristic data that match this parameter value. Enter a complete volume name or press the <Tab> key to complete the name. Wildcard query characters are not supported.

[-iterations <integer>} - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

[-show-flexgroup-as-constituents {true|false}] - Display Flexgroups as Constituents

If the parameter is specified and if the value is true, it will display data for FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

```

cluster1::> qos statistics volume characteristics show -iterations 100
-rows 3
Workload           ID    IOPS      Throughput Request size Read
Concurrency
-----
-----
```

Workload	ID	IOPS	Throughput	Request size	Read
-total-	-	68	176.00KB/s	2650B	7%
8					
vs1vol0-wid102	102	24	96.00KB/s	4096B	20%
13					
vol_1-wid103	103	20	80.00KB/s	4096B	0%
12					
vol_2-wid104	104	1	0KB/s	0B	0%
0					
-total-	-	157	528.00KB/s	3443B	3%
4					
vol_2-wid104	104	48	192.00KB/s	4096B	0%
9					
vol_1-wid103	103	43	172.00KB/s	4096B	0%
0					
vs1vol0-wid102	102	41	164.00KB/s	4096B	14%
6					
-total-	-	274	1016.00KB/s	3797B	2%
2					
vs1vol0-wid102	102	85	340.00KB/s	4096B	8%
4					
vol_2-wid104	104	85	340.00KB/s	4096B	0%
1					
vol_1-wid103	103	84	336.00KB/s	4096B	0%
3					

The example above displays characteristics for the 3 volumes with the highest IOPS and it refreshes the display 100 times before terminating.

```
cluster1::> qos statistics volume characteristics show -vserver vs0  
-volume vs0_volo -iterations 100  


| Workload<br>Concurrency | ID    | IOPS | Throughput | Request Size | Read |
|-------------------------|-------|------|------------|--------------|------|
| -total-<br>2            | -     | 1567 | 783.33KB/s | 512Kb        | 90%  |
| vs0_volo-wid1..<br>1    | 15658 | 785  | 392.33KB/s | 512Kb        | 89%  |
| -total-<br>1            | -     | 1521 | 760.50KB/s | 512Kb        | 90%  |
| vs0_volo-wid1..<br>0    | 15658 | 982  | 491.17KB/s | 512Kb        | 90%  |
| -total-<br>0            | -     | 1482 | 741.00KB/s | 512Kb        | 89%  |
| vs0_volo-wid1..<br>0    | 15658 | 945  | 472.50KB/s | 512Kb        | 90%  |
| -total-<br>0            | -     | 1482 | 741.00KB/s | 512Kb        | 89%  |
| vs0_volo-wid1..<br>0    | 15658 | 945  | 472.50KB/s | 512Kb        | 90%  |
| -total-<br>0            | -     | 1702 | 850.83KB/s | 512Kb        | 90%  |
| vs0_volo-wid1..<br>0    | 15658 | 1018 | 509.00KB/s | 512Kb        | 90%  |


```

The example above displays characteristics for volume `vs0_vo10` in Vserver `vs0` and it refreshes the display 100 times before terminating.

qos statistics volume latency show

Display latency breakdown data per volume

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos statistics volume latency show` command displays the average latencies for volumes on Data ONTAP subsystems.

The command displays the following data:

- The QoS volume name (Workload)
 - The QoS workload ID (ID)
 - Total latency observed per I/O operation (Latency)

- Latency observed per I/O operation in the Network subsystem (Network)
- Latency observed per I/O operation across the internally connected nodes in a Cluster (Cluster)
- Latency observed per I/O operation in the Data management subsystem (Data)
- Latency observed per I/O operation in the Storage subsystem (Disk)
- Latency observed per I/O operation for QoS Policy Group Ceiling (QoS Max)
- Latency observed per I/O operation for QoS Policy Group Floor (QoS Min)
- Latency observed per I/O operation for NVRAM transfer (NVRAM)
- Latency observed per I/O operation for Object Store(Cloud) operations
- Latency observed per I/O operation for FlexCache (FlexCache) operations
- Latency observed per I/O operation for Synchronous Snapmirror (SM Sync) operations
- Latency observed per I/O operation for Volume Activation (VA) operations

The results displayed per iteration are sorted by the total latency field. Each iteration starts with a row that displays the average latency, in microseconds (us) or milliseconds (ms) observed across all volumes.

Parameters

[-node {<nodename>|local}] - Node

Selects the volumes that match this parameter value. If you do not specify this parameter, the command displays data for the entire cluster.

{ [-rows <integer>] - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. The default setting is 10. The allowed range of values is 1 to 20.

| -vserver <vserver name> - Vserver Name

Specifies the Vserver to which the volume belongs.

-volume <volume name> - Volume Name }

Selects the latency data that match this parameter value. Enter a complete volume name or press the <Tab> key to complete the name. Wildcard query characters are not supported.

[-iterations <integer>] - Number of Iterations

Specifies the number of times that the command refreshes the display with updated data before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

[-show-flexport-as-constituents {true|false}] - Display Flexgroups as Constituents

If the parameter is specified and if the value is true, it will display data for FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

cluster1::> qos statistics volume latency show -iterations 100 -rows 3								
Workload		ID	Latency	Network	Cluster	Data	Disk	
Qos	Max	Qos	Min	NVRAM	Cloud	FlexCache	SM Sync	VA
<hr/>								
-total-				110.35ms	110.02ms	0ms	327.00us	0ms
0ms	0ms			0ms	0ms	0ms	0ms	0ms
vs1vol0			111	167.82ms	167.22ms	0ms	603.00us	0ms
0ms	0ms			0ms	0ms	0ms	0ms	0ms
vol1			1234	117.76ms	117.56ms	0ms	191.00us	0ms
0ms	0ms			0ms	0ms	0ms	0ms	0ms
vol2			999	44.24ms	44.05ms	0ms	190.00us	0ms
0ms	0ms			0ms	0ms	0ms	0ms	0ms
-total-			-	38.89ms	38.63ms	0ms	256.00us	0ms
0ms	0ms			0ms	0ms	0ms	0ms	0ms
vol2			999	64.47ms	64.20ms	0ms	266.00us	0ms
0ms	0ms			0ms	0ms	0ms	0ms	0ms
vol1			1234	27.28ms	27.03ms	0ms	253.00us	0ms
0ms	0ms			0ms	0ms	0ms	0ms	0ms
vs1vol0			111	23.72ms	23.47ms	0ms	249.00us	0ms
0ms	0ms			0ms	0ms	0ms	0ms	0ms
-total-			-	409.81ms	409.65ms	0ms	169.00us	0ms
0ms	0ms			0ms	0ms	0ms	0ms	0ms
vol1			1234	816.92ms	816.80ms	0ms	120.00us	0ms
0ms	0ms			0ms	0ms	0ms	0ms	0ms
vol2			999	407.88ms	407.66ms	0ms	219.00us	0ms
0ms	0ms			0ms	0ms	0ms	0ms	0ms
vs1vol0			111	3.68ms	3.49ms	0ms	193.00us	0ms
0ms	0ms			0ms	0ms	0ms	0ms	0ms

The example above displays latencies for the 3 volumes with the highest latencies and it refreshes the display 100 times before terminating.

```

cluster1::> qos statistics volume latency show -vserver vs0 -volume
vs0_vo10 -iterations 100
Workload           ID   Latency     Network    Cluster      Data
Disk    Qos Max   Qos Min    NVRAM       Cloud   FlexCache   SM Sync
VA
-----
----- -----
----- -----
-total-          -  455.00us  158.00us      0ms  297.00us
0ms      0ms      0ms      0ms      0ms      0ms      0ms
0ms
vs0_vo10-wid1.. 15658  428.00us  155.00us      0ms  273.00us
0ms      0ms      0ms      0ms      0ms      0ms      0ms
0ms
-total-          -  337.00us  130.00us      0ms  207.00us
0ms      0ms      0ms      0ms      0ms      0ms      0ms
0ms
vs0_vo10-wid1.. 15658  316.00us  128.00us      0ms  188.00us
0ms      0ms      0ms      0ms      0ms      0ms      0ms
0ms
-total-          -  464.00us  132.00us      0ms  332.00us
0ms      0ms      0ms      0ms      0ms      0ms      0ms
0ms
vs0_vo10-wid1.. 15658  471.00us  130.00us      0ms  341.00us
0ms      0ms      0ms      0ms      0ms      0ms      0ms
0ms
-total-          -  321.00us  138.00us      0ms  183.00us
0ms      0ms      0ms      0ms      0ms      0ms      0ms
0ms
vs0_vo10-wid1.. 15658  302.00us  137.00us      0ms  165.00us
0ms      0ms      0ms      0ms      0ms      0ms      0ms
0ms
-total-          -  418.00us  142.00us      0ms  276.00us
0ms      0ms      0ms      0ms      0ms      0ms      0ms
0ms
vs0_vo10-wid1.. 15658  424.00us  143.00us      0ms  281.00us
0ms      0ms      0ms      0ms      0ms      0ms      0ms
0ms

```

The example above displays latencies for volume `vs0_vo10` in Vserver `vs0` and it refreshes the display 100 times before terminating.

qos statistics volume performance show

Display system performance data per volume

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos statistics volume performance show` command shows the current system performance that each volume is achieving.

The command displays the following data:

- The QoS volume name (Workload)
- The QoS workload ID (ID)
- Input/output operations performed per second (IOPS)
- Throughput in kilobytes per second (KB/s) or megabytes per second (MB/s) as appropriate (Throughput)
- Latency observed per request in microseconds (us) or milliseconds (ms) as appropriate (Latency)

The results displayed per iteration are sorted by IOPS. Each iteration starts with a row that displays the total IOPS used across all volumes. Other columns in this row are either totals or averages.

Parameters

[-node {<nodename>|local}] - Node

Selects the volumes that match this parameter value. If you do not specify this parameter, the command displays data for the entire cluster.

{ [-rows <integer>} - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. The default setting is 10. The allowed range of values is 1 to 20.

| -vserver <vserver name> - Vserver Name

Specifies the Vserver to which the volume belongs.

-volume <volume name> - Volume Name }

Selects the performance data that match this parameter value. Enter a complete volume name or press the `<Tab>` key to complete the name. Wildcard query characters are not supported.

[-iterations <integer>} - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

[-show-flextgroup-as-constituents {true|false}] - Display Flexgroups as Constituents

If the parameter is specified and if the value is true, it will display data for FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

```
cluster1::> qos statistics volume performance show -iterations 100 -rows 3
Workload           ID    IOPS      Throughput     Latency
-----  -----  -----  -----  -----
-total-          -    97      1.90MB/s   216.87ms
vol_2-wid104    104    28      1.75MB/s   412.78ms
vol_1-wid103    103    25      100.00KB/s  169.16ms
vs1vol0-wid102  102    13      52.00KB/s   403.78ms
-total-          -    98      1276.00KB/s 89.98ms
vs1vol0-wid102  102    28      112.00KB/s  80.70ms
vol_1-wid103    103    19      76.00KB/s   114.72ms
vol_2-wid104    104    17      1088.00KB/s 257.60ms
-total-          -    78      1152.00KB/s 225.22ms
vol_1-wid103    103    17      68.00KB/s   452.27ms
vol_2-wid104    104    16      1024.00KB/s 419.93ms
vs1vol0-wid102  102    15      60.00KB/s   210.63ms
```

The example above displays the system performance for the 3 volumes with the highest IOPS and it refreshes the display 100 times before terminating.

```
cluster1::> qos statistics volume performance show -vserver vs0 -volume
vs0_volo -iterations 100
Workload           ID    IOPS      Throughput     Latency
-----  -----  -----  -----  -----
-total-          -    1278     639.17KB/s  404.00us
vs0_volo-wid1..  15658    526     263.17KB/s  436.00us
-total-          -    1315     657.33KB/s  86.00us
vs0_volo-wid1..  15658    528     264.17KB/s  88.00us
-total-          -    1220     609.83KB/s  418.00us
vs0_volo-wid1..  15658    515     257.33KB/s  531.00us
-total-          -    1202     600.83KB/s  815.00us
vs0_volo-wid1..  15658    519     259.67KB/s  924.00us
-total-          -    1240     620.17KB/s  311.00us
vs0_volo-wid1..  15658    525     262.50KB/s  297.00us
```

The example above displays the system performance for volume `vs0_volo` in Vserver `vs0` and it refreshes the display 100 times before terminating.

qos statistics volume resource cpu show

Display CPU resource utilization data per volume

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The qos statistics volume resource cpu show command displays the CPU utilization for volumes per node.

The command displays the following data:

- The QoS volume name (Workload)
- The QoS workload ID (ID)
- CPU utilization observed in percentage (CPU)

The results displayed per iteration are sorted by total CPU utilization. Each iteration starts with a row that displays the total CPU utilization across all volumes.

Parameters

-node {<nodename>|local} - Node

Selects the volumes that match this parameter value.

{ [-rows <integer>} - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. The default setting is 10. The allowed range of values is 1 to 20.

| -vserver <vserver name> - Vserver Name

Specifies the Vserver to which the volume belongs.

-volume <volume name> - Volume Name }

Selects the CPU utilization data that match this parameter value. Enter a complete volume name or press the <Tab> key to complete the name. Wildcard query characters are not supported.

[-iterations <integer>} - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

[-show-flexport-as-constituents {true|false}] - Display Flexgroups as Constituents

If the parameter is specified and if the value is true, it will display data for FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

```

cluster1::> qos statistics volume resource cpu show -node nodeA
-iterations 100 -rows 3
  Workload          ID    CPU
  -----
  --total- (100%)   -    9%
  vs0vol1-wid-102  102   5%
  vs0vol2-wid-121  121   2%
  vs2_volo-wid-..  212   2%
  -total- (100%)   -    8%
  vs0vol1-wid-102  102   5%
  vs0vol2-wid-121  121   2%
  vs2_volo-wid-..  212   1%

```

The example above displays total CPU utilization for the 3 volumes with the highest CPU utilization and it refreshes the display 100 times before terminating.

```

cluster1::> qos statistics volume resource cpu show -node local -vserver
vs0 -volume vs0_voll -iterations 100
  Workload          ID    CPU
  -----
  -total- (100%)   -    2%
  vs0_voll-wid7..  7916  2%
  -total- (100%)   -    2%
  vs0_voll-wid7..  7916  2%
  -total- (100%)   -    1%
  vs0_voll-wid7..  7916  1%
  -total- (100%)   -    2%
  vs0_voll-wid7..  7916  1%
  -total- (100%)   -    2%
  vs0_voll-wid7..  7916  2%

```

The example above displays total CPU utilization for volume `vs0_voll` in Vserver `vs0` and it refreshes the display 100 times before terminating.

qos statistics volume resource disk show

Display disk resource utilization data per volume

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos statistics volume resource disk show` command displays the disk utilization for volumes per node. The disk utilization shows the percentage of time spent on the disk during read and write operations. The command only supports hard disks.

The command displays the following data:

- The QoS volume name (Workload)
- The QoS workload ID (ID)
- Disk utilization (Disk)
- The number of HDD data disks utilized (Number of HDD Disks)

The results displayed are sorted by total disk utilization. Each iteration starts with a row that displays the total disk utilization across all volumes.

Parameters

-node {<nodename>|local} - Node

Selects the volumes that match this parameter value.

{ [-rows <integer>} - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. The default setting is 10. The allowed range of values is 1 to 20.

| -vserver <vserver name> - Vserver Name

Specifies the Vserver to which the volume belongs.

-volume <volume name> - Volume Name }

Selects the disk utilization data that match this parameter value. Enter a complete volume name or press the <Tab> key to complete the name. Wildcard query characters are not supported.

[-iterations <integer>} - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

[-show-flexport-as-constituents {true|false}] - Display Flexgroups as Constituents

If the parameter is specified and if the value is true, it will display data for FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

```

cluster1::> qos statistics volume resource disk show -node nodeB
-iterations 100 -rows 3
      Workload          ID  Disk Number of HDD Disks
----- -----
  -total- (100%)      -   30%           4
  vs0vol1-wid101     101  12%           2
  vs0vol2-wid121     121  10%           1
  vol0-wid1002       1002 8%            1
  -total- (100%)      -   30%           4
  vs0vol1-wid101     101  12%           2
  vs0vol2-wid121     121  10%           1
  vol0-wid1002       1002 8%            1

```

The example above displays total disk utilization for the 3 volumes with the highest disk utilization and it refreshes the display 100 times before terminating.

```

cluster1::> qos statistics volume resource disk show -node local -vserver
vs0 -volume vs0_vo10 -iterations 100
      Workload          ID  Disk Number of HDD Disks
----- -----
  -total-              -   5%           10
  vs0_vo10-wid1..    15658 1%            6
  -total-              -   5%           10
  vs0_vo10-wid1..    15658 1%            6
  -total-              -   6%           10
  vs0_vo10-wid1..    15658 2%            6
  -total-              -   6%           10
  vs0_vo10-wid1..    15658 2%            6
  -total-              -   6%           10
  vs0_vo10-wid1..    15658 2%            6

```

The example above displays total disk utilization for volume `vs0_vo10` in Vserver `vs0` and it refreshes the display 100 times before terminating.

qos statistics workload characteristics show

Display QoS workload characterization

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos statistics workload characteristics show` command displays data that characterizes the behavior of QoS workloads.

The command displays the following data:

- The QoS workload name (Workload)
- The QoS workload ID (ID)
- Input/output operations performed per second (IOPS)
- Throughput achieved in kilobytes per second (KB/s) or megabytes per second (MB/s) as appropriate (Throughput)
- Request size in bytes (B) (Request size)
- Read percentage from total IOPS (Read)
- Concurrency, which indicates the number of concurrent users generating the I/O traffic (Concurrency)

The results displayed per iteration are sorted by IOPS. Each iteration starts with a row that displays the total IOPS used across all QoS workloads. Other columns in this row are either totals or averages.

Parameters

[-node {<nodename>|local}] - Node

Selects the QOS workloads that match this parameter value. If you do not specify this parameter, the command displays data for the entire cluster.

[-iterations <integer>] - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

{ [-rows <integer>] - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. Valid values are from 1 to 20. The default value is 10.

[-policy-group <text>] - QoS Policy Group Name

Selects the QoS workloads that belong to the QoS policy group specified by this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [-adaptive-policy-group <text>] - Adaptive QoS Policy Group Name

Selects the QoS workloads that belong to the QoS adaptive policy group specified by this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [-workload <text>] - QoS Workload Name

Selects the QoS workload that match this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [-workload-id <integer>] - QoS Workload ID }

Selects the QoS workload that match the QoS workload ID specified by this parameter value.

[-show-flexgroup-as-constituents {true|false}] - Display Flexgroups as Constituents

If the parameter is specified and if the value is true, it will display data for FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

```
cluster1::> qos statistics workload characteristics show -iterations 100  
-rows 4  


| Workload<br>Concurrency | ID  | IOPS | Throughput  | Request size | Read |
|-------------------------|-----|------|-------------|--------------|------|
| -total-<br>8            | -   | 68   | 176.00KB/s  | 2650B        | 7%   |
| vs1vol0-wid102<br>13    | 102 | 24   | 96.00KB/s   | 4096B        | 20%  |
| _Scan_Besteff..<br>0    | 101 | 23   | 0KB/s       | 0B           | 0%   |
| vol_1-wid103<br>12      | 103 | 20   | 80.00KB/s   | 4096B        | 0%   |
| vol_2-wid104<br>0       | 104 | 1    | 0KB/s       | 0B           | 0%   |
| -total-<br>4            | -   | 157  | 528.00KB/s  | 3443B        | 3%   |
| vol_2-wid104<br>9       | 104 | 48   | 192.00KB/s  | 4096B        | 0%   |
| vol_1-wid103<br>0       | 103 | 43   | 172.00KB/s  | 4096B        | 0%   |
| vs1vol0-wid102<br>6     | 102 | 41   | 164.00KB/s  | 4096B        | 14%  |
| _Scan_Besteff..<br>0    | 101 | 25   | 0KB/s       | 0B           | 0%   |
| -total-<br>2            | -   | 274  | 1016.00KB/s | 3797B        | 2%   |
| vs1vol0-wid102<br>4     | 102 | 85   | 340.00KB/s  | 4096B        | 8%   |
| vol_2-wid104<br>1       | 104 | 85   | 340.00KB/s  | 4096B        | 0%   |
| vol_1-wid103<br>3       | 103 | 84   | 336.00KB/s  | 4096B        | 0%   |
| _Scan_Besteff..<br>0    | 101 | 20   | 0KB/s       | 0B           | 0%   |


```

The example above displays characteristics for the 4 QoS workloads with the highest IOPS and it refreshes the display 100 times before terminating.

```

cluster1::> qos statistics workload characteristics show -iterations 100
-rows 2 -policy-group pg1
      Workload          ID    IOPS      Throughput Request size Read
Concurrency
-----
-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+-----+-----+
      -total-        -  243   546.86KB/s  2307B  61%
1
0      file-test1_a... 6437    34   136.00KB/s  4096B 100%
0      file-test1_c... 5078    33   133.33KB/s  4096B 100%
0
1      -total-        -  310   3.09MB/s  10428B  55%
0
0      file-test1_a... 6437    36   142.67KB/s  4096B 100%
0      file-test1_b... 9492    35   138.67KB/s  4096B 100%
0
1      -total-        -  192   575.71KB/s  3075B  71%
0
0      file-test1-wi... 7872    39   157.33KB/s  4096B 100%
0      file-test1_c... 5078    38   153.33KB/s  4096B 100%
0

```

The example above displays the characteristics for the 2 QoS workloads belonging to QoS policy group *pg1* with the highest IOPS and it refreshes the display 100 times before terminating.

```

cluster1::> qos statistics workload characteristics show -iterations 100
-wkload-id 9492
      Workload          ID    IOPS      Throughput Request size Read
Concurrency
-----
-----+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+-----+
      -total-        -    737     2.14MB/s   3045B  79%
1
0      file-test1_b... 9492    265     1058.67KB/s  4096B 100%
1
1      -total-        -    717     4.26MB/s   6235B  80%
1
0      file-test1_b... 9492    272     1086.67KB/s  4096B 100%
1
0      -total-        -    623     2.50MB/s   4202B  86%
0
0      file-test1_b... 9492    263     1050.67KB/s  4096B 100%
0
0      -total-        -    595     2.11MB/s   3712B  89%
0
0      file-test1_b... 9492    266     1064.00KB/s  4096B 100%
0

```

The example above displays the characteristics for the QoS workload with QoS workload ID 9492 and it refreshes the display 100 times before terminating.

qos statistics workload latency show

Display latency breakdown data per QoS workload

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos statistics workload latency show` command displays the average latencies for QoS workloads on Data ONTAP subsystems.

The command displays the following data:

- The QoS workload name (Workload)
- The QoS workload ID (ID)
- Total latency observed per I/O operation (Latency)
- Latency observed per I/O operation in the Network subsystem (Network)
- Latency observed per I/O operation across the internally connected nodes in a Cluster (Cluster)
- Latency observed per I/O operation in the Data management subsystem (Data)

- Latency observed per I/O operation in the Storage subsystem (Disk)
- Latency observed per I/O operation for QoS Policy Group Ceiling (QoS Max)
- Latency observed per I/O operation for QoS Policy Group Floor (QoS Min)
- Latency observed per I/O operation for NVRAM transfer (NVRAM)
- Latency observed per I/O operation for Object Store(Cloud) operations
- Latency observed per I/O operation for FlexCache (FlexCache) operations
- Latency observed per I/O operation for Synchronous Snapmirror (SM Sync) operations
- Latency observed per I/O operation for Volume Activation (VA) operations

The results displayed per iteration are sorted by the total latency field. Each iteration starts with a row that displays the average latency, in microseconds (us) or milliseconds (ms) observed across all QoS workloads.

Parameters

[`-node {<nodename>|local}`] - Node

Selects the QoS workloads that match this parameter value. If you do not specify this parameter, the command displays data for the entire cluster.

[`-iterations <integer>`] - Number of Iterations

Specifies the number of times that the command refreshes the display with updated data before terminating. If you do not specify this parameter, the command continues to run until you interrupt it by pressing Ctrl-C.

[`-refresh-display {true|false}`] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

{ [`-rows <integer>`] - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. Valid values are from 1 to 20. The default value is 10.

[`-policy-group <text>`] - QoS Policy Group Name

Selects the QoS workloads that belong to the QoS policy group specified by this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [`-adaptive-policy-group <text>`] - Adaptive QoS Policy Group Name

Selects the QoS workloads that belong to the QoS adaptive policy group specified by this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [`-workload <text>`] - QoS Workload Name

Selects the QoS workload that match this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [`-workload-id <integer>`] - QoS Workload ID }

Selects the QoS workload that match the QoS workload ID specified by this parameter value.

[`-show-flexgroup-as-constituents {true|false}`] - Display Flexgroups as Constituents

If the parameter is specified and if the value is true, it will display data for FlexVols and Flexgroup

Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

cluster1::> qos statistics workload latency show -iterations 100 -rows 3							
Workload	ID	Latency	Network	Cluster	Data	Disk	
Qos	Max	Qos Min	NVRAM	Cloud	FlexCache	SM Sync	VA
-total-		110.35ms	110.02ms	0ms	327.00us	0ms	
0ms	0ms	0ms	0ms	0ms	0ms	0ms	
vs1vol0	111	167.82ms	167.22ms	0ms	603.00us	0ms	
0ms	0ms	0ms	0ms	0ms	0ms	0ms	
vol1	1234	117.76ms	117.56ms	0ms	191.00us	0ms	
0ms	0ms	0ms	0ms	0ms	0ms	0ms	
vol2	999	44.24ms	44.05ms	0ms	190.00us	0ms	
0ms	0ms	0ms	0ms	0ms	0ms	0ms	
-total-	-	38.89ms	38.63ms	0ms	256.00us	0ms	
0ms	0ms	0ms	0ms	0ms	0ms	0ms	
vol2	999	64.47ms	64.20ms	0ms	266.00us	0ms	
0ms	0ms	0ms	0ms	0ms	0ms	0ms	
vol1	1234	27.28ms	27.03ms	0ms	253.00us	0ms	
0ms	0ms	0ms	0ms	0ms	0ms	0ms	
vs1vol0	111	23.72ms	23.47ms	0ms	249.00us	0ms	
0ms	0ms	0ms	0ms	0ms	0ms	0ms	
-total-	-	409.81ms	409.65ms	0ms	169.00us	0ms	
0ms	0ms	0ms	0ms	0ms	0ms	0ms	
vol1	1234	816.92ms	816.80ms	0ms	120.00us	0ms	
0ms	0ms	0ms	0ms	0ms	0ms	0ms	
vol2	999	407.88ms	407.66ms	0ms	219.00us	0ms	
0ms	0ms	0ms	0ms	0ms	0ms	0ms	
vs1vol0	111	3.68ms	3.49ms	0ms	193.00us	0ms	
0ms	0ms	0ms	0ms	0ms	0ms	0ms	

The example above displays latencies for the 3 QoS workloads with the highest latencies and it refreshes the display 100 times before terminating.

```

cluster1::> qos statistics workload latency show -iterations 100 -rows 2
-policy-group pg1
Workload           ID   Latency     Network    Cluster      Data
Disk   Qos Max   Qos Min   NVRAM       Cloud   FlexCache   SM Sync
VA
-----
----- -----
----- -----
-total-          -   4.80ms  287.00us   0ms  427.00us
4.08ms        0ms   0ms        0ms        0ms   0ms        0ms
0ms
file-test1-wi.. 7872  9.60ms  265.00us   0ms  479.00us
8.85ms        0ms   0ms        0ms        0ms   0ms        0ms
0ms
file-test1_a... 6437  8.22ms  262.00us   0ms  424.00us
7.53ms        0ms   0ms        0ms        0ms   0ms        0ms
0ms
-total-          -   4.20ms  296.00us   0ms  421.00us
3.48ms        0ms   0ms        0ms        0ms   0ms        0ms
0ms
file-test1-wi.. 7872  8.70ms  211.00us   0ms  489.00us
8.00ms        0ms   0ms        0ms        0ms   0ms        0ms
0ms
file-test1_a... 6437  6.70ms  297.00us   0ms  464.00us
5.94ms        0ms   0ms        0ms        0ms   0ms        0ms
0ms
-total-          -   5.90ms  303.00us   0ms  1.71ms
3.88ms        0ms   0ms        0ms        0ms   0ms        0ms
0ms
file-test1-wi.. 7872  11.36ms 263.00us   0ms  2.06ms
9.04ms        0ms   0ms        0ms        0ms   0ms        0ms
0ms
file-test1_a... 6437  9.48ms  250.00us   0ms  2.30ms
6.93ms        0ms   0ms        0ms        0ms   0ms        0ms
0ms

```

The example above displays latencies for the 2 QoS workloads belonging to QoS policy group *pg1* with the highest IOPS and it refreshes the display 100 times before terminating.

The example above displays the latencies for the QoS workload with QoS workload ID 9492 and it refreshes the display 100 times before terminating.

qos statistics workload performance show

Display system performance data per QoS workload

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `qos statistics workload performance show` command shows the current system performance that each QoS workload is achieving.

The command displays the following data:

- The QoS workload name (Workload)
- The QoS workload ID (ID)
- Input/output operations performed per second (IOPS)
- Throughput in kilobytes per second (KB/s) or megabytes per second (MB/s) as appropriate (Throughput)
- Latency observed per request in microseconds (us) or milliseconds (ms) as appropriate (Latency)

The results displayed per iteration are sorted by IOPS. Each iteration starts with a row that displays the total IOPS used across all QoS workloads. Other columns in this row are either totals or averages.

Parameters

[-node {<nodename>|local}] - Node

Selects the QoS workloads that match this parameter value. If you do not specify this parameter, the command displays data for the entire cluster.

[-iterations <integer>] - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

{ [-rows <integer>] - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. Valid values are from 1 to 20. The default value is 10.

[-policy-group <text>] - QoS Policy Group Name

Selects the QoS workloads that belong to the QoS policy group specified by this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [-adaptive-policy-group <text>] - Adaptive QoS Policy Group Name

Selects the QoS workloads that belong to the QoS adaptive policy group specified by this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [-workload <text>] - QoS Workload Name

Selects the QoS workload that match this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [-workload-id <integer>] - QoS Workload ID }

Selects the QoS workload that match the QoS workload ID specified by this parameter value.

[-show-flextgroup-as-constituents {true|false}] - Display Flexgroups as Constituents

If the parameter is specified and if the value is true, it will display data for FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

```
cluster1::> qos statistics workload performance show -iterations 100 -rows 4
Workload          ID    IOPS      Throughput     Latency
-----  -----
-total-           -    97       1.90MB/s   216.87ms
_Scan_Besteff..  101   31        0KB/s     0ms
vol_2-wid104    104   28       1.75MB/s   412.78ms
vol_1-wid103    103   25      100.00KB/s  169.16ms
vs1vol0-wid102  102   13      52.00KB/s   403.78ms
-total-           -    98      1276.00KB/s 89.98ms
_Scan_Besteff..  101   34        0KB/s     0ms
vs1vol0-wid102  102   28      112.00KB/s  80.70ms
vol_1-wid103    103   19      76.00KB/s   114.72ms
vol_2-wid104    104   17      1088.00KB/s 257.60ms
-total-           -    78      1152.00KB/s 225.22ms
_Scan_Besteff..  101   30        0KB/s     0ms
vol_1-wid103    103   17      68.00KB/s   452.27ms
vol_2-wid104    104   16      1024.00KB/s 419.93ms
vs1vol0-wid102  102   15      60.00KB/s   210.63ms
```

The example above displays the system performance for the 4 QoS workloads with the highest IOPS and it refreshes the display 100 times before terminating.

```
cluster1::> qos statistics workload performance show -iterations 100 -rows 2 -policy-group pg1
Workload          ID    IOPS      Throughput     Latency
-----  -----
-total-           -    2598     9.96MB/s  1223.00us
file-testfile..  4228   650      2.54MB/s  1322.00us
file-testfile..  11201   635      2.48MB/s  1128.00us
-total-           -    2825     10.89MB/s 714.00us
file-testfile..  4228   707      2.76MB/s  759.00us
file-testfile..  11201   697      2.72MB/s  693.00us
-total-           -    2696     10.13MB/s 1149.00us
file-testfile..  4228   645      2.52MB/s  945.00us
file-testfile..  6827   634      2.48MB/s  1115.00us
```

The example above displays the system performance for the 2 QoS workloads belonging to QoS policy group pg1 with the highest IOPS and it refreshes the display 100 times before terminating.

```

cluster1::> qos statistics workload performance show -iterations 100
-workload-id 11201
Workload          ID    IOPS      Throughput     Latency
-----  -----  -----  -----  -----
-total-          -    2866    10.92MB/s   905.00us
file-testfile..  11201    674     2.63MB/s   889.00us
-total-          -    2761    10.55MB/s  1054.00us
file-testfile..  11201    638     2.49MB/s  1055.00us
-total-          -    2810    10.58MB/s  832.00us
file-testfile..  11201    685     2.68MB/s  909.00us
-total-          -    2593    9.86MB/s  1092.00us
file-testfile..  11201    632     2.47MB/s  964.00us

```

The example above displays the system performance for the QoS workload with QoS workload ID `11201` and it refreshes the display `100` times before terminating.

qos statistics workload resource cpu show

Display CPU resource utilization data per QoS workload

Availability: This command is available to *cluster* administrators at the `admin` privilege level.

Description

The `qos statistics workload resource cpu show` command displays the CPU utilization for QoS workloads per node.

The command displays the following data:

- The QoS workload name (Workload)
- The QoS workload ID (ID)
- CPU utilization observed in percentage (CPU)

The results displayed per iteration are sorted by total CPU utilization. Each iteration starts with a row that displays the total CPU utilization across all QoS workloads.

Parameters

-node {<nodename>|local} - Node

Selects the QOS workloads that match this parameter value.

[-iterations <integer>] - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

{ [-rows <integer>] - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. Valid values are from 1 to 20. The default value is 10.

[-policy-group <text>] - QoS Policy Group Name

Selects the QoS workloads that belong to the QoS policy group specified by this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [-adaptive-policy-group <text>] - Adaptive QoS Policy Group Name

Selects the QoS workloads that belong to the QoS adaptive policy group specified by this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [-workload <text>] - QoS Workload Name

Selects the QoS workload that match this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [-workload-id <integer>] - QoS Workload ID }

Selects the QoS workload that match the QoS workload ID specified by this parameter value.

[-show-flexgroup-as-constituents {true|false}] - Display Flexgroups as Constituents

If the parameter is specified and if the value is true, it will display data for FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

```
cluster1::> qos statistics workload resource cpu show -node nodeA  
-iterations 100 -rows 3  
Workload          ID    CPU  
-----  
--total- (100%)   -    9%  
vs0-wid-102      102   5%  
file-bigvmdk...   121   2%  
vs2_volo-wid...  212   2%  
-total- (100%)   -    8%  
vs0-wid-101      102   5%  
file-bigvmdk...   121   2%  
vs2_volo-wid...  212   1%
```

The example above displays total CPU utilization for the 3 QoS workloads with the highest CPU utilization and it refreshes the display 100 times before terminating.

```
cluster1::> qos statistics workload resource cpu show -node local
-iterations 100 -rows 2 -policy-group pg1
Workload          ID    CPU
-----
-total- (100%)   -    41%
file-test1_b...  9492  16%
file-test1_c...  5078  16%
-total- (100%)   -    43%
file-test1_c...  5078  17%
file-test1_b...  9492  16%
-total- (100%)   -    40%
file-test1_c...  5078  16%
file-test1_b...  9492  15%
```

The example above displays total CPU utilization for the 2 QoS workloads belonging to QoS policy group *pg1* with the highest IOPS and it refreshes the display 100 times before terminating.

```
cluster1::> qos statistics workload resource cpu show -node local
-iterations 100 -workload-id 9492
Workload          ID    CPU
-----
-total- (100%)   -    15%
file-test1_b...  9492  3%
-total- (100%)   -    14%
file-test1_b...  9492  3%
-total- (100%)   -    14%
file-test1_b...  9492  2%
-total- (100%)   -    13%
file-test1_b...  9492  3%
```

The example above displays total CPU utilization for the QoS workload with QoS workload ID 9492 and it refreshes the display 100 times before terminating.

qos statistics workload resource disk show

Display disk resource utilization data per QoS workload

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The *qos statistics workload resource disk show* command displays the disk utilization for QoS workloads per node. The disk utilization shows the percentage of time spent on the disk during read and write operations. The command displays disk utilization for system-defined workloads; however, their disk utilization is not included in the total utilization. The command only supports hard disks.

The command displays the following data:

- The QoS workload name (Workload)
- The QoS workload ID (ID)
- Disk utilization (Disk)
- The number of HDD data disks utilized (Number of HDD Disks)

The results displayed are sorted by total disk utilization. Each iteration starts with a row that displays the total disk utilization across all QoS workloads.

Parameters

-node {<nodename>|local} - Node

Selects the QoS workloads that match this parameter value.

[-iterations <integer>] - Number of Iterations

Specifies the number of times the display is refreshed before terminating. If you do not specify this parameter, the command iterates until interrupted by Ctrl-C.

[-refresh-display {true|false}] - Toggle Screen Refresh Between Each Iteration

Specifies the display style. If true, the command clears the display after each data iteration. If false, the command displays each data iteration below the previous one. The default is false.

{ [-rows <integer>] - Number of Rows in the Output

Specifies the number of busiest QoS policy groups to display. Valid values are from 1 to 20. The default value is 10.

[-policy-group <text>] - QoS Policy Group Name

Selects the QoS workloads that belong to the QoS policy group specified by this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [-adaptive-policy-group <text>] - Adaptive QoS Policy Group Name

Selects the QoS workloads that belong to the QoS adaptive policy group specified by this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [-workload <text>] - QoS Workload Name

Selects the QoS workload that match this parameter value. If you do not specify this parameter, the command displays data for all QoS workloads.

| [-workload-id <integer>] - QoS Workload ID }

Selects the QoS workload that match the QoS workload ID specified by this parameter value.

[-show-flexgroup-as-constituents {true|false}] - Display Flexgroups as Constituents

If the parameter is specified and if the value is true, it will display data for FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

```

cluster1::> qos statistics workload resource disk show -node nodeB
-iterations 100 -rows 3
      Workload          ID  Disk Number of HDD Disks
----- -----
      -total- (100%)    -   30%           4
      _RAID            -   20%           4
      vs0-wid101       101  12%           2
      file-1-wid121    121  10%           1
      vol0-wid1002     1002 8%            1
      _WAFL            -   7%            3
      -total- (100%)    -   30%           4
      vs0-wid101       101  12%           2
      file-1-wid121    121  10%           1
      _RAID            -   10%           4
      vol0-wid1002     1002 8%            1
      _WAFL            -   7%            3

```

The example above displays total disk utilization for the 3 QoS workloads with the highest disk utilization and it refreshes the display 100 times before terminating.

```

cluster1::> qos statistics workload resource disk show -node local
-iterations 100 -rows 2 -policy-group pg1
      Workload          ID  Disk Number of HDD Disks
----- -----
      -total-           -   3%           10
      file-test1_a...   6437 6%            6
      file-test1_wi...  7872 6%            6
      -total-           -   3%           10
      file-test1_a...   6437 5%            6
      file-test1_wi...  7872 5%            6
      -total-           -   3%           10
      file-test1_a...   6437 6%            6
      file-test1_wi...  7872 6%            6

```

The example above displays total disk utilization for the 2 QoS workloads belonging to QoS policy group *pg1* with the highest IOPS and it refreshes the display 100 times before terminating.

```

cluster1::> qos statistics workload resource disk show -node local
-iterations 100 -workload-id 6437
Workload           ID  Disk Number of HDD Disks
-----
-total-          -   3%        10
file-test1_a...  6437  6%        6
-total-          -   3%        10
file-test1_a...  6437  5%        6
-total-          -   3%        10
file-test1_a...  6437  6%        6

```

The example above displays total disk utilization for the QoS workload with QoS workload ID 6437 and it refreshes the display 100 times before terminating.

qos workload commands

qos workload delete

Delete workload

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

Deletes a workload from a cluster. When you delete a workload, the associated data object is no longer controlled by its policy group.

You can only delete user-defined workloads. You cannot delete system-defined or preset workloads.

Parameters

-workload <text> - Workload Name

Specifies the name of the workload that you want to delete.

Examples

```

cluster1::> qos workload delete workload1

```

Deletes the "workload1" user-defined workload from the "cluster1" cluster.

qos workload show

Display a list of workloads

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

Shows the current status of workloads on a cluster. Use this command to determine the types of workloads that are currently on a cluster. The types of workloads include: system-defined, preset, and user-defined. The system generates system-defined and preset workloads. You cannot create, modify, or delete these workloads. Also, you can only modify or delete a user-defined workload , but cannot create one

Parameters

{ [-fields <fieldname>, ...]

If you specify the -fields <fieldname>, ... parameter, the command output also includes the specified field or fields. You can use '-fields ?' to display the fields to specify.

| [-instance] }

If you specify the -instance parameter, the command displays detailed information about all fields.

[-workload <text>] - Workload Name

If you use this parameter, the command displays the workloads that contain the specified workload name.

[-uuid <UUID>] - Workload UUID (privilege: advanced)

If you use this parameter, the command displays the workloads that contain the specified UUID.

[-class <QoS Configuration Class>] - Workload Class

If you use this parameter, the command displays the workloads that contain the specified class. The Class options include system-defined, preset, and user-defined.

[-wid <integer>] - Workload ID

If you use this parameter, the command displays the workloads that contain the specified internal workload ID.

[-category <text>] - Workload Category

If you use this parameter, the command displays the workloads that contain the specified category. The category options include Scanner and Efficiency.

[-policy-group <text>] - Policy Group Name

If you use this parameter, the command displays the workloads that match the specified policy group name.

[-read-ahead <text>] - Read-ahead Tunables

If you use this parameter, the command displays the workloads that contain the specified read-ahead cache tunable.

[-vserver <vserver name>] - Vserver

If you use this parameter, the command displays the workloads that match the specified Vserver.

[-volume <volume name>] - Volume

If you use this parameter, the command displays the workloads that match the specified volume.

[-qtree <qtree name>] - Qtree Name

If you use this parameter, the command displays the workloads that match the specified Qtree name.

[-lun <text>] - LUN Name

If you use this parameter, the command displays the workloads that match the specified LUN name.

[-file <text>] - File Path

If you use this parameter, the command displays the workloads that match the specified file path.

[-max-throughput <qos_tput>] - Maximum Throughput

Selects the workloads that match this parameter value

A maximum throughput limit specifies the throughput in IOPS that the workload must not exceed.

[-min-throughput <qos_tput>] - Minimum Throughput

Selects the workloads that match this parameter value

A minimum throughput specifies the desired performance level for a workload in IOPS.

[-is-adaptive {true|false}] - Adaptive

If you use this parameter, the command displays only adaptive workloads.

[-is-constituent {true|false}] - Is Constituent Volume

If this parameter is specified, the command displays information only about storage objects that either are or are not constituents of a FlexGroup, depending on the value provided.

Examples

```
cluster1::> qos workload show -class user-defined
Workload      Wid   Policy Group Vserver  Volume    LUN     Qtree   File
Path
-----
-----
vs2-wid100    100   pg1          vs2      -        -        -        -
```

Shows all user-defined workloads and the corresponding storage objects on the "cluster1" cluster.

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—with 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.