



qos statistics commands

ONTAP 9.10.1 commands

NetApp
February 11, 2024

Table of Contents

qos statistics commands	1
qos statistics characteristics show	1
qos statistics latency show	3
qos statistics performance show	6
qos statistics resource cpu show	8
qos statistics resource disk show	10
qos statistics volume characteristics show	12
qos statistics volume latency show	15
qos statistics volume performance show	18
qos statistics volume resource cpu show	20
qos statistics volume resource disk show	22
qos statistics workload characteristics show	24
qos statistics workload latency show	28
qos statistics workload performance show	32
qos statistics workload resource cpu show	35
qos statistics workload resource disk show	37

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-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 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.

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-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 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		
vslvol0	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		
vslvol0	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		
vslvol0	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-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 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-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 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.

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.