storage automated-working-set-analyzer commands

ONTAP 9.10.1 commands

NetApp
September 27, 2022

Table of Contents

storage automated-working-set-analyzer commands

storage automated-working-set-analyzer show

Display running instances

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

Description

The automated-working-set-analyzer show command displays the Automated Working-set Analyzer running instances.

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.

[-node <nodename>] - Node Name

This parameter indicates the node name that the AWA instance runs on.

[-flash-cache {true|false}] - Flash Cache Node-wide Modeling

This parameter indicates whether the AWA is modeling flash-cache.

[-aggregate-uuid <UUID>] - Uuid of the Aggregate

This parameter indicates the aggregate uuid that the AWA instance runs on.

[-aggregate <aggregate name>] - Aggregate

This parameter indicates the aggregate name that the AWA instance runs on.

[-working-set-size {true|false}] - Working Set Size

This parameter indicates whether the AWA instance is configured to find the working set size.

[-start-time <Date>] - Starting Time

This parameter indicates the time when the AWA instance was started.

[-total-intervals <integer>] - Total Interval Count

This parameter indicates the total number of intervals that the AWA instance has covered.

[-read-throughput {<integer>[Bps|KBps|MBps|GBps]}] - Read Throughput

This parameter indicates the maximum read throughput over an interval that AWA has observed from the storage disks.

[-write-throughput {<integer>[Bps|KBps|MBps|GBps]}] - Write Throughput

This parameter indicates the maximum write throughput over an interval that AWA has observed to the storage disks

[-cacheable-read <percent>] - Cacheable Read

This parameter indicates the maximum percent of cacheable read over an interval that AWA has observed. Cacheable reads are non-sequential reads, i.e., the percentage of data reads that could have been cached.

[-cacheable-write <percent>] - Cacheable Write

This parameter indicates the maximum percent of cacheable write over an interval that AWA has observed. Cacheable writes are random overwrites, percentage of disk writes that could have been cached.

[-projected-cache-size {<integer>[KB|MB|GB|TB|PB]}] - Max Projected Cache Size

This parameter indicates the projected Flash Pool cache usage.

[-projected-read-hit <percent>] - Projected Read Hit

This parameter indicates the percentage of blocks that could be read from the Flash Pool cache instead of HDDs.

[-projected-write-hit <percent>] - Projected Write Hit

This parameter indicates the percentage of block overwrites that could go to the Flash Pool cache instead of HDDs.

[-referenced-interval-id <integer>] - Referenced Interval ID

This parameter indicates the interval in which the cache size effect information is derived from.

[-referenced-interval-time <Date>] - Referenced Interval Time

This parameter indicates the time when the referenced interval for the cache size effect information is derived from.

[-referenced-interval-cache-size {<integer>[KB|MB|GB|TB|PB]}] - Referenced Interval Cache Size

This parameter indicates the cache size at the end of the referenced interval from which the cache size effect information is based on.

[-read-hit-20 <percent>] - 20% Cache Read Hit

This parameter indicates the predicted read hit rate when the cache size is 20% of the referenced cache size.

[-read-hit-40 <percent>] - 40% Cache Read Hit

This parameter indicates the predicted read hit rate when the cache size is 40% of the referenced cache size.

[-read-hit-60 <percent>] - 60% Cache Read Hit

This parameter indicates the predicted read hit rate when the cache size is 60% of the referenced cache size.

[-read-hit-80 <percent>] - 80% Cache Read Hit

This parameter indicates the predicted read hit rate when the cache size is 80% of the referenced cache size.

[-read-hit-100 <percent>] - 100% Cache Read Hit

This parameter indicates the predicted read hit rate when the cache size is 100% of the referenced cache size.

[-write-hit-20 <percent>] - 20% Cache Write Hit

This parameter indicates the predicted write hit rate when the cache size is 20% of the referenced cache size.

[-write-hit-40 <percent>] - 40% Cache Write Hit

This parameter indicates the predicted writehit rate when the cache size is 40% of the referenced cache size.

[-write-hit-60 <percent>] - 60% Cache Write Hit

This parameter indicates the predicted write hit rate when the cache size is 60% of the referenced cache size.

[-write-hit-80 <percent>] - 80% Cache Write Hit

This parameter indicates the predicted write hit rate when the cache size is 80% of the referenced cache size.

[-write-hit-100 <percent>] - 100% Cache Write Hit

This parameter indicates the predicted write hit rate when the cache size is 100% of the referenced cache size.

[-num-intervals-show <integer>] - Number of intervals to show

This parameter indicates the number of intervals to the past this command is showing.

Examples

The following example shows a running instance of automated-working-set-analyzer on node *node1* for aggregate *aggr0* .

```
cluster1::> cluster-1::*> storage automated-working-set-analyzer show
Node           FC    Aggregate    wss    Intervals Start Time
-------------- ----- ------------ ------ ---------
------------------------
node1          false aggr0        false        125 Wed Jul 22 13:58:17
2015
```

storage automated-working-set-analyzer start

Command to start Automated Working Set Analyzer on node or aggregate

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

Description

The automated-working-set-analyzer start command enables the Automated Workload Analyzer that is capable of doing the following:

- Flash Pool modeling for an aggregate
- Flash Cache modeling for a node - can not specify an aggregate.
- Working set size estimation
- Workload monitoring

Parameters

`-node <nodename>` - Node Name

This parameter indicates the node name that the AWA instance runs on.

`[-flash-cache {true|false}]` - Flash Cache Node-wide Modeling

This parameter indicates whether the AWA is modeling flash-cache.

`[-aggregate <aggregate name>]` - Aggregate

This parameter indicates the aggregate name that the AWA instance runs on.

`[-working-set-size {true|false}]` - Working Set Size

This parameter indicates whether the AWA instance is configured to find the working set size.

Examples

```
cluster1::> storage automated-working-set-analyzer start -node vsim1
-aggregate aggr0
```

storage automated-working-set-analyzer stop

Command to stop Automated Working Set Analyzer on node or aggregate

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

Description

The storage automated-working-set-analyzer stop command terminates one or multiple Automated Workload Analyzer running instances.

Parameters

`-node <nodename>` - Node Name

This parameter indicates the node name that the AWA instance runs on.

`[-flash-cache {true|false}]` - Flash cache node-wide modeling

This parameter indicates whether the AWA is modeling flash-cache.

[-aggregate <aggregate name>] - Aggregate

This parameter indicates the aggregate name that the AWA instance runs on.

Examples

```
cluster1::> storage automated-working-set-analyzer stop -node vsim1
-aggregate aggr1
```

storage automated-working-set-analyzer volume show

Displays the Automated Working Set Analyzer volume table

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

Description

The automated-working-set-analyzer volume show command displays the volume statistics reported by the corresponding Automated Working-set Analyzer running instances.

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.

[-node <nodename>] - Node

This parameter indicates the node name that the AWA instance runs on.

[-flash-cache {true|false}] - Flash Cache Node-wide Modeling

This parameter indicates whether the AWA is modeling flash-cache.

[-vol-uuid <UUID>] - Uuid of the Volume

This parameter indicates the volume uuid that this command is issued on.

[-aggregate <aggregate name>] - Aggregate

This parameter indicates the aggregate name that the AWA instance runs on.

[-volume <volume name>] - Volume

This parameter indicates the volume name that this command is issued on.

[-rank <integer>] - Cache Benefit Rank

This parameter indicates the rank of this volume among all volumes that would be most benefited by the modeled cache technology based on the AWA prediction.

[-read-throughput {<integer>[Bps|KBps|MBps|GBps]}] - Read Throughput

This parameter indicates the maximum read throughput over an interval that AWA has observed from the storage disks for this volume.

[-write-throughput {<integer>[Bps|KBps|MBps|GBps]}] - Write Throughput

This parameter indicates the maximum write throughput over an interval that AWA has observed to the storage disks for this volume.

[-cacheable-read <percent>] - Cacheable Read

This parameter indicates the maximum percent of cacheable read over an interval that AWA has observed for this volume. Cacheable reads are non-sequential reads, i.e., the percentage of data reads that could have been cached.

[-cacheable-write <percent>] - Cacheable Write

This parameter indicates the maximum percent of cacheable write over an interval that AWA has observed. Cacheable writes are random overwrites, percentage of disk writes that could have been cached.

[-projected-cache-size {<integer>[KB|MB|GB|TB|PB]}] - Max Projected Cache Size

This parameter indicates the projected Flash Pool cache usage by this volume.

[-projected-read-hit <percent>] - Projected Read Hit

This parameter indicates the percentage of blocks that could be read from the Flash Pool cache instead of HDDs for this volume.

[-projected-write-hit <percent>] - Projected Write Hit

This parameter indicates the percentage of block overwrites that could go to the Flash Pool cache instead of HDDs for this volume.

[-num-intervals-show <integer>] - Number of intervals to show

This parameter indicates the number of intervals to the past this command is showing.

Examples

```
cluster1::> cluster-1::*> storage automated-working-set-analyzer volume
show
Node           FC    Aggregate    Volume       Rank  Read Thrupt Write
Thrupt
-------------- ----- ------------ ------------ ---- ------------
------------
vsim1          false aggr0        vol0            1   230.47KBps
580.09KBps
```