



storage automated-working-set-analyzer commands

Command reference

NetApp
July 17, 2026

Table of Contents

- storage automated-working-set-analyzer commands 1
 - storage automated-working-set-analyzer show 1
 - Description 1
 - Parameters 1
 - Examples 3
 - storage automated-working-set-analyzer start 3
 - Description 4
 - Parameters 4
 - Examples 4
 - storage automated-working-set-analyzer stop 4
 - Description 4
 - Parameters 4
 - Examples 5
 - storage automated-working-set-analyzer volume show 5
 - Description 5
 - Parameters 5
 - Examples 6

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 ] (privilege: advanced) }`

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

`[-node <nodename>] - Node Name (privilege: advanced)`

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

`[-flash-cache {true|false}] - Flash Cache Node-wide Modeling (privilege: advanced)`

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

`[-aggregate-uuid <UUID>] - Uuid of the Aggregate (privilege: advanced)`

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

`[-aggregate <aggregate name>] - Aggregate (privilege: advanced)`

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

`[-working-set-size {true|false}] - Working Set Size (privilege: advanced)`

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

`[-start-time <Date>] - Starting Time (privilege: advanced)`

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

`[-total-intervals <integer>] - Total Interval Count (privilege: advanced)`

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

`[-read-throughput {<integer>[Bps|KBps|MBps|GBps]}] - Read Throughput (privilege: advanced)`

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 (privilege: advanced)

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

[-cacheable-read <percent>] - Cacheable Read (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

This parameter indicates the projected Flash Pool cache usage.

[-projected-read-hit <percent>] - Projected Read Hit (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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

[-referenced-interval-time <Date>] - Referenced Interval Time (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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

[-flash-cache {true|false}] - Flash Cache Node-wide Modeling (privilege: advanced)

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

[-aggregate <aggregate name>] - Aggregate (privilege: advanced)

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

[-working-set-size {true|false}] - Working Set Size (privilege: advanced)

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 (privilege: advanced)

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

[-flash-cache {true|false}] - Flash cache node-wide modeling (privilege: advanced)

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

[`-aggregate <aggregate name>`] - Aggregate (privilege: advanced)

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`](privilege: advanced) }

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

[`-node <nodename>`] - Node (privilege: advanced)

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

[`-flash-cache {true|false}`] - Flash Cache Node-wide Modeling (privilege: advanced)

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

[`-vol-uuid <UUID>`] - Uuid of the Volume (privilege: advanced)

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

[`-aggregate <aggregate name>`] - Aggregate (privilege: advanced)

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

[`-volume <volume name>`] - Volume (privilege: advanced)

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

[`-rank <integer>`] - Cache Benefit Rank (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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

[-projected-read-hit <percent>] - Projected Read Hit (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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
```


Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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